Stockout Response Optimization Engine: A Tool for Strategic Recovery Planning

by

Scott Marintsch
Bachelor of Business Administration, Texas A&M University, 2011

and

Sharada Maruti Prabhu
Master of Business Administration, Indian Institute of Management Lucknow, 2020

SUBMITTED TO THE PROGRAM IN SUPPLY CHAIN MANAGEMENT
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF APPLIED SCIENCE IN SUPPLY CHAIN MANAGEMENT
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2025

Signature of Author: ______________________________
Department of Supply Chain Management
May 09, 2025

Signature of Author: ______________________________
Department of Supply Chain Management
May 09, 2025

Certified by: ______________________________
Dr. Milena Janjevic
Director of MIT Supply Chain Design Lab
Capstone Advisor

Accepted by: ______________________________
Prof. Yossi Sheffi
Director, Center for Transportation and Logistics
Elisha Gray II Professor of Engineering Systems
Professor, Civil and Environmental Engineering

Stockout Response Optimization Engine: A Tool for Strategic Recovery Planning

by

Scott Marintsch
Bachelor of Business Administration, Texas A&M University, 2011

and

Sharada Maruti Prabhu
Master of Business Administration, Indian Institute of Management Lucknow, 2020

SUBMITTED TO THE PROGRAM IN SUPPLY CHAIN MANAGEMENT
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF APPLIED SCIENCE IN SUPPLY CHAIN MANAGEMENT
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2025

Signature of Author: ____
Department of Supply Chain Management
May 09, 2025

Signature of Author: ____
Department of Supply Chain Management
May 09, 2025

Certified by: ____
Dr. Milena Janjevic
Director of MIT Supply Chain Design Lab
Capstone Advisor

Accepted by: ____
Prof. Yossi Sheffi
Director, Center for Transportation and Logistics
Elisha Gray II Professor of Engineering Systems
Professor, Civil and Environmental Engineering

Stockout Response Optimization Engine: A Tool for Strategic Recovery Planning

by

Scott Marintsch

and

Sharada Maruti Prabhu

Submitted to the Program in Supply Chain Management
on May 09, 2025, in Partial Fulfillment of the
Requirements for the Degree of Master of Applied Science in Supply Chain Management

ABSTRACT

Stockouts in hub-and-spoke distribution networks can disrupt service levels and drive-up logistics costs. This project develops a Mixed-Integer Linear Programming (MILP) model to support inventory recovery decisions by integrating forecasted stockout risk, inventory positions, demand, and operational constraints. The model evaluates hub-to-DC and lateral DC-to-DC transfers, balancing trade-offs between transportation cost, underutilized trucks, early transfer penalties, and service failure costs. A key feature is its configurability at the SKU and customer level, allowing planners to simulate tailored recovery actions. The model consumes externally provided stockout probability forecasts at the distribution center (DC) level to inform recovery prioritization. Results show the model can reduce stockouts to zero under feasible conditions while highlighting the cost trade-offs of recovery options. The model is deployed through an interactive dashboard to support scenario analysis and rapid planning. This framework provides supply chain planners with a data-driven approach to proactively mitigate stockouts and improve network resilience.

Capstone Advisor: Dr. Milena Janjevic

Title: Director of MIT Supply Chain Design Lab

ACKNOWLEDGMENTS

We are grateful to Dr. Milena Janjevic for her guidance, insight, and steady support throughout this project. Her mentorship helped shape our thinking and approach at every stage.

We also thank Armada for the opportunity to engage with a real-world challenge, and for their collaboration, transparency, and valuable input throughout the process.

We appreciate the resources, structure, and learning environment provided by MIT, which made this experience both rigorous and rewarding.

Finally, we thank our families and peers for their continued encouragement and support.

TABLE OF CONTENTS

1 INTRODUCTION

1.1 Motivation

Historically, Third-Party Logistics (3PL) providers have primarily focused on core activities for their customers, serving as day-to-day operators by managing warehouses, booking truckloads, and acting as executors within their clients' supply chains. However, since the 2010s, 3PLs have evolved from simple service providers to strategic partners, delivering value-added services that have paved the way for the emergence of the 4PL model. These expanded services include managing customer data and developing software tools, enabling customers to gain deeper insights and greater visibility into their own supply chains through enhanced collaboration with their partners.

Supply chain visibility and orchestration technology, which entails coordinating, managing, and optimizing all supply chain activities, are becoming increasingly vital as recent logistical disruptions, such as the COVID-19 pandemic, have highlighted the high cost of unmitigated risks in end-to-end supply networks. Today, supply chains face unprecedented pressure from shifting consumer expectations, frequent disruptions, and operational challenges like labor shortages and capacity constraints. External factors such as weather events, cyberattacks, and port closures continue to expose vulnerabilities, making visibility and orchestration technology indispensable for anticipating disruptions and maintaining smooth operations.

Armada, our capstone sponsor, is a company that focuses on the execution and optimization of end-to-end supply chains for their customers, primarily within the food and beverage vertical. They service large quick-service restaurant (QSR) clients, such as McDonald’s and Chik-fil-A, and manage freight programs utilizing a 4PL type business model. In addition to distribution warehousing and transportation services, the company maintains a continued focus on enhancing digital experiences and technology services for their customers.

Armada is driven to prevent stockouts and late shipments by developing proactive alerting technology powered by machine learning for their customers. Their goal is to utilize prescriptive tools that monitor critical milestone events and identify risks of missing commitments or Service Level Agreements (SLAs). Ultimately, Armada seeks to scale these tools to enhance their existing warehousing and transportation services, facilitating seamless change management and promoting end-user adoption of advanced technologies.

1.2 Problem Statement

Armada operates a multi-tier hub-and-spoke distribution network consisting of centralized hubs, downstream distribution centers (DCs), and retail clients. Inventory flows from hubs to DCs on a weekly basis, and DCs fulfill customer demand. Each DC can be connected to one or multiple hubs, and demand uncertainty at the DC level leads to frequent stockout risks.

To mitigate these risks, Armada currently uses an in-house predictive model that estimates the weekly probability of stockouts over a four-week horizon for each hub-SKU pair. These probabilities are generated based on SKU inventory, forecasted demand, inbound supply, and transit data, and are computed at the hub level. However, the current recovery process is entirely manual: supply chain planners review these predictions and rely on domain knowledge to decide whether and how to respond (e.g., transfer inventory between hubs or reroute truckloads).

This reactive and manual approach leads to inconsistent responses, suboptimal truck utilization, and missed opportunities to mitigate downstream stockouts. Armada seeks to move toward a more automated and prescriptive solution that converts network-wide stockout forecasts into proactive recovery recommendations.

Armada currently uses an in-house stockout model to predict the probability of stockouts for each SKU as it moves from a hub to a distribution center over a four-week period. Planners use the model's outputs to address inventory shortages, manually determining recovery actions based on their judgment. The company seeks to transition to a more automated, prescriptive recovery process that leverages network-wide stockout predictions to generate optimal recovery strategy recommendations. Armada aims to scale its response optimization engine to enhance risk interpretation and seamlessly integrate prescriptive analytics into the workflows of supply chain practitioners and planners.

In this context, the capstone addresses the following questions:

1. How can a recommendation engine help end users develop optimized recovery strategies that mitigate stockouts and delays, while balancing transfer costs and service level risk across the network?
2. How can inventory transfers be optimized within the hub–DC network to reduce empty truck space and minimize transportation usage, while accounting for forecasted stockout risk and SKU-level demand?
3. When should inventory optimization levers—such as safety stock repositioning or preemptive inventory balancing—be applied to minimize the need for short-term recovery actions?

The recovery strategies under consideration include:

- Expediting existing hub-to-DC shipments when timely fulfillment is at risk
- Executing lateral DC-to-DC transfers to shift available stock to higher-risk locations
- Leveraging existing scheduled truck capacity ("free rides") to reduce incremental transportation cost
- Routing shipments through cross-dock facilities to consolidate loads and improve truck utilization
- Delaying action or accepting a stockout when the cost of recovery exceeds the expected penalty

1.3 Scope: Project Goals and Expected Outcomes

This project develops a recommendation engine that generates automated recovery strategies in response to potential disruptions within Armada's supply chain. By leveraging network data and stockout model outputs, the engine generates actionable recommendations to mitigate stockouts, reduce delays, and optimize inventory flow across the network.

The engine leverages Armada's stockout predictions and network data as key inputs to generate prescriptive recovery strategies. By analyzing these inputs, the engine assists planners in making data-driven decisions to address stockouts, optimize inventory transfers within the hub and DC network, and enhance overall transfer efficiency. It focuses on reducing empty truck space and enhancing the flow of inventory, while also providing a foundation for future integration of inventory optimization methods, such as safety stock management and inventory balancing, to further optimize supply chain performance and minimize unnecessary transfers.

The deliverables to Armada include:

- A proof of concept featuring a backend optimization model and a lightweight front-end interface, demonstrating the engine's ability to translate stockout predictions into actionable recovery strategies.
- An analysis of potential courses of action based on disruptions and suggesting the optimal response by weighing factors such as time, cost and inventory.

This initiative aims to incorporate prescriptive analytics into Armada's supply chain decision-making process, enabling more efficient and automated responses to disruptions and reducing reliance on manual interventions. Additionally, based on the research and development from this effort, Armada aims to build a digital playbook which can be scaled and executed for future innovation in this area.

2 STATE OF THE PRACTICE

Stockouts remain a persistent challenge in modern supply chains, especially within multi-tier distribution networks that rely on timely product movement across hubs, distribution centers, and retail endpoints. With rising demand variability, transportation delays, and capacity constraints, organizations are under increasing pressure to maintain service levels while minimizing inventory and logistics costs. To manage this complexity, companies adopt a combination of preventive and mitigative strategies—each targeting different stages of the disruption lifecycle.

In this section, we begin by reviewing the common causes of stockouts (Section 2.1), followed by a summary of the main response strategies observed in both industry and literature. While a comprehensive review of all possible strategies is beyond the scope of this chapter, the most pertinent response strategies have been examined.

These strategies fall into two primary categories:

- **Prevention** – Strategies designed to reduce the likelihood of stockouts before they occur by building resilience into inventory and planning systems. Example strategies include:
    a. Dynamic Inventory Policies (Section 2.2)
    b. Multi-Echelon Inventory Optimization (Section 2.3)
    c. Risk Pooling (Section 2.4)
- **Mitigation** – Strategies that respond after a stockout risk is detected, aiming to minimize service disruption through operational interventions. Example strategies include:
    a. Route optimization (Section 2.5)
    b. Inventory Redistribution (Section 2.6)

Supply chain visibility plays a foundational role in both categories by enabling real-time data access, early warning signals, and coordinated decision-making (Section 2.7). This capstone project aligns with the mitigation stream, with a focus on two specific recovery strategies:

- Hub-to-DC Replenishment Optimization, which improves shipment timing and efficiency from hubs to DCs; and
- A Proposed Inventory Redistribution Policy, which uses a Mixed-Integer Linear Programming (MILP) model to optimize lateral DC-to-DC transfers for balancing inventory and minimizing stockouts across the network.

2.1 Stockout Management

Armada primarily serves the Quick Service Restaurant (QSR) industry, which is growing rapidly as brands embrace new technologies, respond to shifting consumer demands, and increasingly focus on operational efficiency to serve more customers, reduce costs, and scale faster. In this evolving environment, effective inventory management is essential for optimizing operations, controlling costs, and ensuring customer satisfaction. Recent trends show a decline in inventory levels as a percentage of sales, with stocks being placed closer to end customers (Silver et al., 2017). To keep pace, companies are adopting innovative strategies to streamline inventory management.

Stockouts are a frequent challenge in supply chains and can be particularly expensive for organizations. Defined as the occurrence when available stock depletes to zero (Silver et al., 2017), stockouts can stem from various factors, such as unexpected surges in demand or inadequate inventory levels, including safety stock, to fulfill orders (see Figure 1). The consequences of stockouts are far-reaching: they lead to lost sales, heightened customer attrition risk, decreased customer satisfaction, and a gradual erosion of brand loyalty (Silver et al., 2017).

Figure 1. Primary causes of stockouts in the restaurant industry (Nigam, 2016)

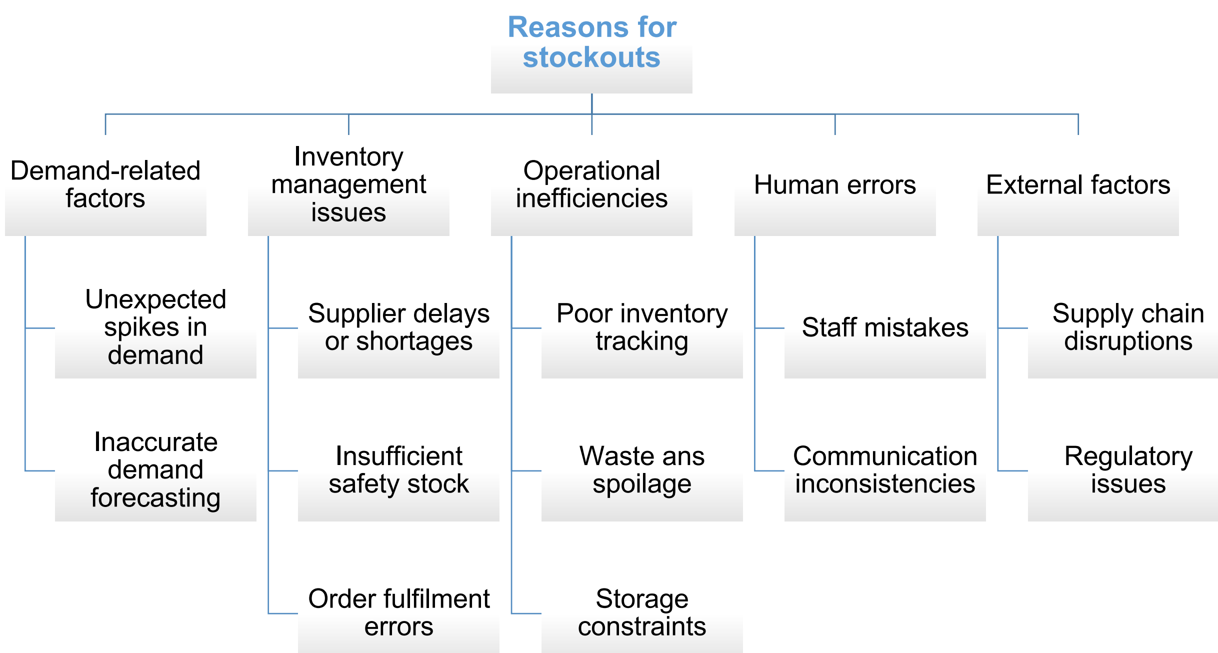

Our literature review explores strategies employed by researchers and companies across various industries to reduce the risk of stockouts. Conventional approaches to minimizing stockouts include Safety Stock Optimization, Multi-Echelon Inventory Optimization (MEIO), Risk Pooling, Route Optimization, and Inventory Redistribution, and Dynamic Inventory policies.

2.2 Dynamic Inventory Policies

Dynamic inventory policies offer a structured approach to mitigating supply and demand variability by tailoring replenishment decisions to operational contexts. These policies form the foundation of proactive stockout prevention, enabling firms to maintain desired service levels while controlling inventory-related costs.

2.2.1 Replenishment and Review-Based Policies

Replenishment policies determine the timing and quantity of inventory orders and serve as the foundation for dynamic inventory control. The five common policies are reviewed below and summarized in Figure 2.

Figure 2. Types of Dynamic Inventory Policies (Silver, et al., 2017)

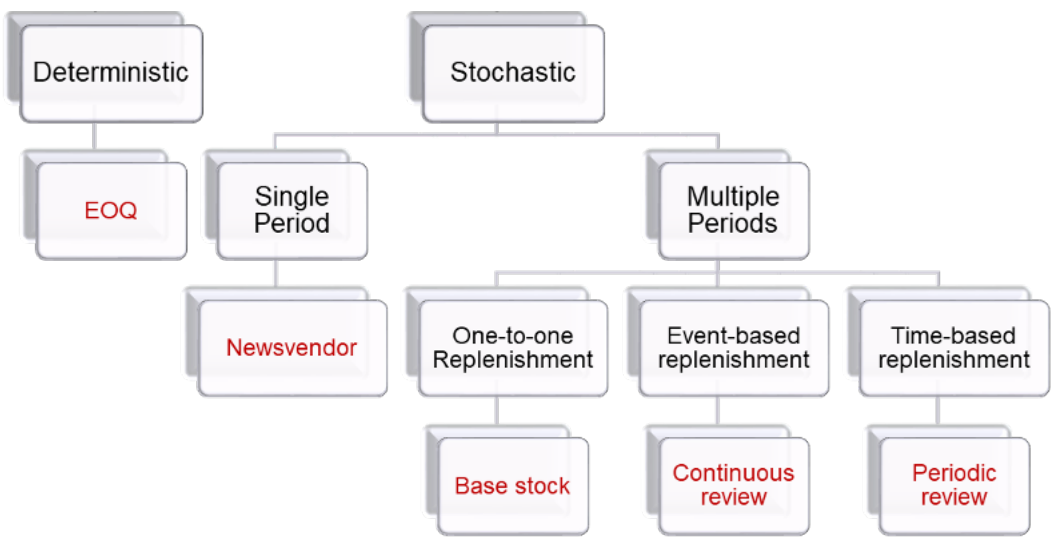

Economic Order Quantity (EOQ) optimizes holding and ordering costs for stable, predictable demand, ensuring timely stock replenishment. The **Single Period/Newsvendor Model** addresses the risks of overstocking and understocking in situations with finite-horizon, uncertain demand, making it especially suitable for perishable or seasonal products (Caplice & Ponce, 2024). Cao et al. (2018) enhanced the newsvendor model by incorporating decision-making under risk, showing that considering factors like loss aversion, buybacks, and stockout penalties can optimize order quantities and effectively reduce stockouts.

For continuous operations, the Base Stock Policy ensures timely replenishment, maintaining high service levels in settings such as e-commerce or just-in-time systems. The **Continuous Review Policy (s, Q)** constantly monitors inventory and triggers fixed-quantity orders when stock drops below a reorder point, effectively managing demand fluctuations. Meanwhile, the **Periodic Review Policy (R, S)** reviews inventory at regular intervals and replenishes it to a target level, making it ideal for environments where continuous tracking is not feasible (Caplice & Ponce, 2024).

2.2.2 Safety Stock Optimization

Safety Stock Optimization (SSO) refers to the process of determining the optimal level of extra inventory a company should maintain to protect against uncertainties in demand or supply. Safety stock acts as a buffer to minimize stockouts, enabling businesses to maintain service levels while avoiding excessive inventory holding costs.

The below safety stock formula accounts for demand and lead time variability to prevent stockouts. It calculates the buffer inventory required to meet a desired service level using the service level factor (Z), lead time (LT), average demand (D), and the variances of demand (σ_D^2) and lead time (σ_{LT}^2). The formula combines these factors to ensure adequate inventory in uncertain conditions.

$$Safety\ Stock = Z \times \sqrt{(LT \times \sigma_D^2) + (D^2 \times \sigma_{LT}^2)}$$

- Z: Safety factor based on the desired service level.
- σ_D^2: Variance in demand.
- σ_{LT}^2: Variance in lead time.
- LT: Average lead time.
- D: Average demand.

To manage demand uncertainty, four safety stock models can be considered (Brunaud et al., 2018):

1. Proportional to throughput
2. Proportional to throughput with risk-pooling effects

3. Explicit risk-pooling
4. Guaranteed service time

These models enable the simultaneous optimization of safety stock, reserve stock, and base stock levels while integrating material flow considerations into supply chain planning. Technologies and approaches such as machine learning and simulation modeling can be utilized to uncover patterns, analyze interdependencies, account for seasonality, and leverage real-time data for more accurate and adaptive supply chain planning.

2.3 Multi-Echelon Inventory Optimization (MEIO)

A common approach to reducing stockout events is utilizing a strategic technique called Multi-Echelon Inventory Optimization (MEIO) to optimize inventory across multiple nodes within a network. MEIO considers the interdependencies between the different tiers of a supply chain, such as hubs and DCs, and considers balancing inventory levels to optimize holding costs and service levels while minimizing total costs and ensuring product availability.

Historically, organizations have concentrated on managing inventories at individual sites, often leading to inefficiencies and surplus stock across the broader supply chain (Driessen, 2023). MEIO provides a solution by considering the entire supply chain, optimizing inventory levels at each stage. This comprehensive approach reduces unnecessary safety stock and improves overall supply chain efficiency while ensuring that customer service levels are maintained (Driessen, 2023).

Zhang et al. (2021) analyzed a supply chain with five enterprise nodes, demonstrating that a centralized MEIO strategy led to lower overall costs while maintaining high customer service levels. Their findings show that MEIO strategies can optimize total inventory across a supply chain by effectively allocating inventory among echelons based on demand, cost, and lead time data. This balancing act is crucial for minimizing inventory holding costs while ensuring that demand is met. As part of MEIO, replenishment policies, such as order-up-to levels, order quantities, and base stock levels, define when and how much inventory to order at each echelon to maintain service levels while minimizing costs. These policies are vital for balancing inventory throughout the supply chain, helping to avoid stockouts and reduce excess inventory (Axsäter, 2006).

2.4 Risk Pooling

Risk pooling involves consolidating inventory across fewer locations to reduce demand and lead time variability, helping minimize stockouts and lower the overall need for safety stock. By sharing demand forecasts, sales data, and inventory levels, supply chain stakeholders can respond more swiftly to shortages, leading to more stable and efficient inventory management (Oeser, 2015). The primary benefit of risk pooling is its ability to enhance demand forecasting accuracy and align inventory levels more closely with actual demand, improving service levels without requiring excess stock at each location. The net benefit is the reduction of total expected holding and penalty costs in an inventory system through the consolidation of demand (Eppen, 1979).

However, a potential downside of risk pooling is the increased transportation cost that may result from consolidating inventory in fewer locations, potentially necessitating more frequent shipments to satisfy local demand (Sharma & Bojorquez Aispuro, 2020). Additionally, concentrating inventory in fewer locations can expose the supply chain to risks, as delays or issues at a central hub could disrupt operations (Sharma & Bojorquez Aispuro, 2020). Despite these challenges, when implemented correctly, risk pooling can be a powerful strategy for reducing uncertainty and optimizing inventory management across the network.

2.5 Route Optimization

An Inventory Routing Problem (IRP) is a logistical optimization challenge that integrates inventory management, determining when and how much to restock to avoid stockouts or overstocking, with routing optimization, which focuses on planning the most efficient delivery routes to minimize transportation costs while fulfilling inventory requirements. Soysal, et al. (2015) highlight that a key challenge in applying IRP models to the food sector is the limitation imposed by the short shelf life of products.

Moin and Salhi (2007) classified solutions to the Inventory Routing Problem (IRP) into two primary approaches:

1. **Theoretical Approach**: This approach focuses on determining lower bounds for the problem by dividing it into two phases: inventory management and routing (e.g., solving the traveling salesman problem). These phases are addressed iteratively until an optimal solution is found.
2. **Practical Approach**: This approach seeks near-optimal solutions using heuristics and mathematical models such as Integer Linear Programming (ILP) or Mixed-Integer Programming (MIP). These techniques are used to optimize specific aspects, including delivery schedules, quantities, inventory levels, or route selection.

To minimize stockouts and improve supply chain efficiency for Armada, optimizing routes from hubs to distribution centers (DCs) and between DCs is essential. A systematic approach includes:

1. **Hub-to-DC Replenishment Optimization:** Ensuring timely and efficient replenishment of inventory at DCs based on forecasted downstream DC and restaurant demand.
2. **DC-to-DC Transfers**: Optimizing routes for redistributing inventory between DCs to address demand surges or mitigate stockouts.

Key inputs for Armada's route optimization include:

- **Inventory Data**: Stock levels, reorder points, and safety stock thresholds at DCs.
- **Demand Forecasts**: Consumption patterns at the DC and potentially the restaurant level.
- **Route Information**: Distance, travel time, and road conditions between hubs and DCs, as well as among DCs.
- **Constraints**: Delivery windows, vehicle capacity, perishability, and shelf life of goods.

2.6 Inventory Redistribution Policy

In a hub and distribution center (DC) network, Bertrand and Bookbinder's (1998) redistribution policies help minimize total inventory costs by reallocating stock between the hub and DCs to balance the marginal costs of holding inventory and stockouts across the network. Marginal holding costs refer to the expenses of storing additional inventory at the hub or DC, such as storage, handling, and insurance, which increase as inventory levels rise. In contrast, marginal stockout costs occur when inventory is insufficient at a hub or DC to meet demand, leading to lost sales, customer dissatisfaction, and potentially expensive expedited shipments to prevent further stockouts.

A redistribution policy focused on reducing future stockout events ensures that these marginal costs are balanced between the hub and DCs before the next replenishment cycle. The use of lateral transshipments, defined as the redistribution of inventory between locations within the same level or tier of a supply chain, to address demand imbalances (Bertrand and Bookbinder, 1998), can significantly mitigate the likelihood of stockouts within a supply chain network (see Figure 3).

Figure 3. Hub-and-Spoke Distribution Network with DC-to-DC Transfers (i.e., Lateral Transshipments)

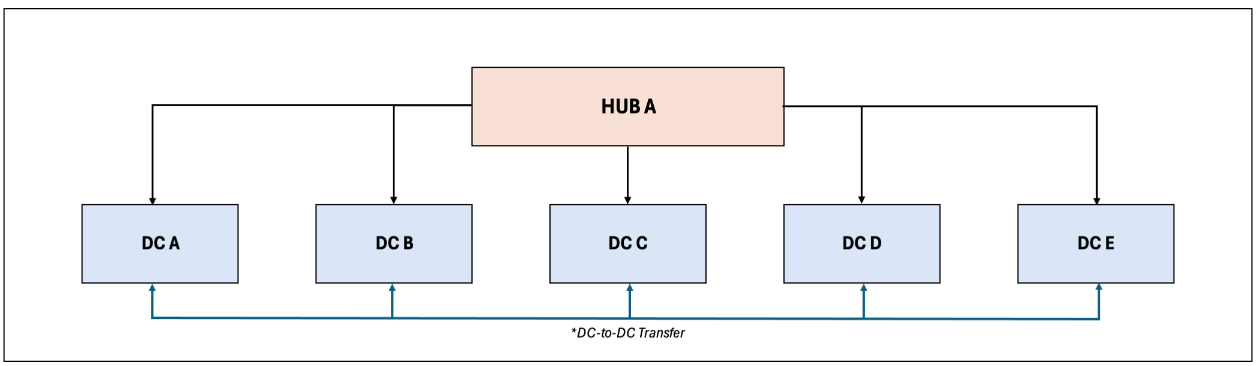

A lateral transshipment policy, known as service level adjustment (SLA), was proposed by Lee et al. (2006), where both emergency lateral transshipment (ELT) and preventive lateral transshipment (PLT) were integrated as part of a unified transshipment strategy within a supply chain network. These policies were shown to improve customer satisfaction, reduce stockout events, and lower overall transfer costs within the network.

Within the framework of Inventory Redistribution Policy (IRP), Halbe and Je (2022) reviewed the use of Mixed-Integer Linear Programming (MILP) to design inventory rebalancing strategies through lateral transshipments, optimizing inventory transfers across multi-echelon supply chain networks. This model utilized heuristics and inventory classification systems to prioritize high-impact SKUs when determining optimal transfer strategies. Their findings revealed that implementing lateral transshipment could reduce total inventory costs by 10% to 25%, while also improving inventory positions across a medical device company's network.

2.7 Supply Chain Visibility

Improving supply chain visibility between hubs and Distribution Centers (DCs) helps reduce stockout events by enabling real-time data sharing, enhancing inventory management, and aligning replenishment strategies. Tools like collaborative planning, forecasting, and replenishment (CPFR) facilitate early identification of potential stockouts, enabling quicker response times and better risk management (Hill et al., 2018). Enhanced end-to-end visibility improves customer service and fosters a more efficient, agile supply chain.

Panahifar et al. (2015) further advanced the understanding of CPFR by proposing a comprehensive framework that identifies and categorizes key factors, such as CPFR enablers, barriers, and partner selection criteria, helping businesses navigate the challenges of effective collaboration across supply chain partners. Additionally, Adoolffo et al. (2023) demonstrate how integrating Industry 4.0 technologies, including IoT, big data, and cloud computing, into CPFR models enhances forecasting accuracy and decision-making, providing a roadmap for using these technologies to strengthen supply chain collaboration and reduce stockouts.

3 METHODOLOGY

This chapter outlines the methodology used to design a decision-support model that proactively mitigates stockout risks in a hub-and-spoke distribution network. The approach builds upon Armada's existing in-house stockout model and addresses current limitations in how recovery decisions are made. Armada currently uses a predictive model that estimates stockout probabilities at the **hub–SKU level**, based on forecasted demand, inventory, and inbound shipments. However, this model has two key limitations:

- It does not provide DC-specific predictions, limiting granularity for localized decision-making
- Recovery decisions based on the model are manual and reactive, with no formal framework to weigh cost trade-offs across possible strategies

The goal of this capstone is to build a prescriptive framework that:

- Refines and localizes stockout predictions at the distribution center (DC) level, and
- Recommends optimal recovery strategies based on network-wide constraints and cost trade-offs

The methodology is divided into two core components:

1. **Model Development and Risk Quantification:** This section details the foundational assumptions, input data sources, and preprocessing steps necessary for constructing a robust model aimed at accurately interpreting and quantifying stockout risks at the distribution center (DC) level
2. **Mixed-Integer Linear Programming (MILP) Optimization Model:** This section introduces a MILP optimization model that utilizes identified risk indicators, in conjunction with cost considerations, capacity constraints, and scheduling parameters. This model generates prescriptive recovery actions, such as replenishments, transfers, and delay strategies

The methodology selection was motivated by current industry practices and supported by relevant literature, ensuring alignment with established approaches to supply chain risk management. The chapter progresses from data preparation and estimation techniques to the formulation of input assumptions critical to model accuracy. The chapter concludes with a detailed delineation of the structure of the optimization model. This includes an examination of the model's objective function, pivotal constraints, and the configurability of key parameters through a user-centric interface.

Our methodology provides a systematic framework for enhancing operational resilience within a hub-and-spoke distribution network, contributing to the proactive mitigation of stockout risks and fostering efficient inventory management practices.

3.1 Overview of Model Inputs, Preparation Approach, and Execution Process

This section provides a high-level overview of the model inputs, preparation steps, and overall execution process that guided the development of the stockout response optimization engine. It serves as a roadmap for the detailed methodology discussions that follow. Figure 4 summarizes the project's core workflow, from input preparation through model-driven recommendation generation.

Figure 4. Core Workflow: Model Preparation to Execution

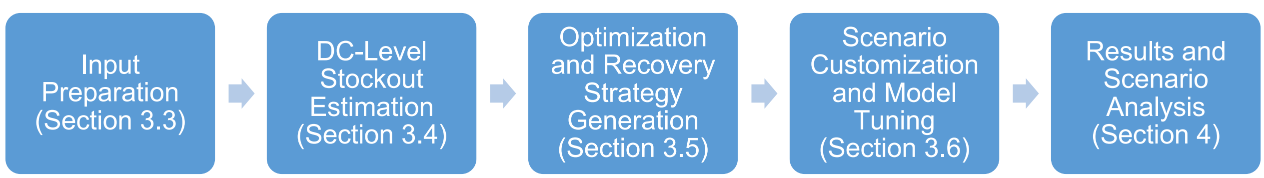

The project leveraged Armada's operational datasets as its foundation, including item master records, facility attributes, transportation cost matrices, existing truck schedules, stockout risk demand files, and on-hand inventory positions. These core inputs are described in detail in Section 3.3 and Appendix B. Given that Armada's stockout probability forecasts were originally generated at the hub level, a manual extrapolation process was employed to estimate distribution center (DC)-level stockout risks. This estimation method, based on historical demand allocations and inventory coverage assessments, is outlined in Section 3.4.

Following stockout risk estimation, a Mixed-Integer Linear Programming (MILP) model was developed to prescribe recovery strategies. The model integrates transportation costs, truck underutilization penalties, early transfer penalties, and stockout penalties to recommend replenishment, transfer, or delay actions. The full optimization model formulation and decision logic are detailed in Section 3.5 and Appendix A. Ultimately, the outputs of the modelling process are deployed through a dashboard interface that allows planners to explore recovery scenarios by adjusting penalty weights and operational parameters (Section 3.6 and Section 4.2).

3.2 Methodology Selection

Based on the literature reviewed and the current practices of Armada's supply chain planning teams, we identified two key mechanisms to enhance resilience and optimize responses to potential stockouts:

- **Hub-to-DC Replenishment Optimization:** This strategy optimizes hub-to-DC replenishments to minimize stockouts, ensure efficient transportation, and balance inventory during periods of heightened demand.
- **Proposed Inventory Redistribution Policy**: Utilizing a MILP optimization model, Armada seeks to optimize DC-to-DC transfers through lateral transshipment execution. This approach enhances the efficiency of inventory redistribution across the network.

Armada currently aggregates demand planning and forecasting data from each hub and DC within its network. Leveraging historical data and disruption patterns, they have developed a stockout prediction model that forecasts stockout probabilities per SKU at the hub level. These weekly forecasts, spanning up to four weeks in advance, serve as critical inputs alongside supply and demand data and current SKU stock positions for the MILP model.

The MILP model operationalizes the aforementioned strategies by determining optimal hub-to-DC replenishment plans to address projected shortfalls and triggering DC-to-DC lateral transfers when redistribution is more efficient. This approach enhances response times and improves inventory availability network-wide.

The optimal responses to potential stockout events have been refined into the core **Orchestration Recovery Strategies**, which are decision points generated by the MILP model at the individual product level based on predicted stockouts, supply constraints, and cost-efficiency:

1. **Expediting Purchase Orders**: Evaluating the need to expedite existing hub-to-DC Purchase Orders (POs) or initiate new POs to meet downstream DC demand.
2. **DC-to-DC Transfer**: Assessing the viability of lateral transshipment from neighboring DCs to mitigate stockout risks.
3. **Leveraging Scheduled Truck Capacity ("Free Rides"):** Utilizing uncommitted capacity on pre-scheduled transportation routes to reposition at-risk inventory without incurring incremental transportation cost. This approach transforms underutilized freight capacity into a low-cost recovery mechanism.
4. **Cross-Docking Shipments**: Routing product flows through cross-dock facilities to consolidate partial loads, enhance trailer utilization, and improve the efficiency of downstream delivery. This strategy is particularly effective when direct full-load shipments are operationally viable but cost-prohibitive due to low utilization or fragmented demand.
5. **Delay the Recovery Event or Accept the Stockout:** Postponing recovery actions to allow for more informed risk and cost assessments. In cases where recovery costs outweigh the impact of a stockout, the model may recommend accepting the stockout as a cost-effective option.

Additional considerations integrated into the MILP model ensure that decisions are aligned with product and customer value:

- **Customer Priority**: Prioritization of key customer accounts using a cost model that reflects the strategic value of each end-customer. This ensures that stockout risks are not shifted from lower-priority customers to high-value strategic ones, where the impact of stockouts is more significant.
- **SKU Importance**: Inclusion of a SKU Classification System that evaluates the strategic impact of SKUs on inventory and demand planning, ensuring that key

products, such as high-velocity or seasonal/promotional items, are prioritized in stockout mitigation efforts.

This systematic approach benefits from its application across Armada's entire supply chain, optimizing decision-making beyond individual DC or SKU levels. Building on the project's core development of an Optimization Engine, the system integrates the MILP model to evaluate stockout risks systematically over the upcoming four weeks, identifying optimal recovery strategies on a network-wide scale prior to each replenishment cycle. An essential outcome of this optimization effort is the reduction of unused truck space and minimization of truck numbers through consolidated inventory transfers between hubs and DCs. By leveraging predictive stockout data, Armada aims to enhance logistics efficiency significantly, moving away from current ad-hoc transfer methods that often result in underutilized trailer capacity.

This evaluation of recovery strategies involves rigorous assessment of associated risks and costs, including the evaluation of alternative options such as DC-to-DC inventory transfers or accelerated hub-based POs. This analytical approach ensures the selection of the most cost-effective and risk-averse strategies to mitigate stockouts and maintain operational continuity.

3.3 Key Assumptions, Data and Data Pre-Processing Requirements

To support optimization, our model requires refined stockout predictions and a set of high-quality operational inputs. While Armada's current in-house model estimates stockout probabilities at the hub–SKU level, it does not offer DC-specific risk predictions and does not explicitly incorporate inventory availability or localized demand exposure. In contrast, our approach builds a DC-level stockout estimation framework that:

- Allocates hub-level risk signals to downstream DCs using historical demand splits
- Explicitly considers **on-hand and in-transit inventory** at each DC
- Evaluates whether each DC has sufficient coverage for projected demand
- Generates **granular, forward-looking risk probabilities** that serve as inputs to the recovery optimization engine

These refined predictions allow for **localized, cost-aware decision-making** that the current hub-level model does not support. The effectiveness of the optimization model is contingent upon several critical assumptions and the availability of high-quality data. The model's performance relies on rigorous data pre-processing to ensure that all inputs are both accurate and consistent. Key assumptions and data requirements include:

1. **Stockout Predictions:** The model relies on a pre-existing stockout prediction system that furnishes detailed forecasts at the distribution center (DC) level. These predictions encompass SKU-specific details such as quantities and stockout horizons, crucial for setting inventory constraints and determining replenishment schedules.
2. **On-Hand Inventory Positions:** Accurate, real-time visibility into on-hand inventory across the network is pivotal, particularly for SKUs susceptible to imminent stockouts. This data forms the basis of inventory constraints, ensuring that replenishment strategies align with actual network conditions. It encompasses current inventory levels, outbound allocations (to-be-picked), and visibility into in-transit or inbound shipments (to-be-filled).
3. **Existing Routes and Schedules:** Availability of scheduled transportation routes for the current week is assumed. This dataset encompasses both Hub-to-DC and DC-to-DC routes, including details on temperature-controlled trucks (Dry, Cooler, Frozen). This information is essential for identifying optimal transshipment opportunities and planning logistics effectively.
4. **Transportation Costs and Spot Rates:** Comprehensive data on transportation costs, including fixed and variable components, is essential. Incorporating spot rate information ensures the model reflects real-time cost fluctuations accurately. Armada's provided spot rate data informs projected transportation costs, crucial for optimizing new truckload provisions.
5. **Item Master Data and Priority Classification:** Detailed SKU attributes from the item master, such as temperature classification, case quantities, and cases per pallet, are necessary. These attributes facilitate precise shipment planning aligned with temperature-controlled logistics and capacity constraints. SKU classification

aids in prioritizing stockouts based on strategic importance, influencing penalty weighting in cost calculations.

6. **Facility Data and Priority Classification:** Facility data, including facility classification priority, supports strategic replenishment prioritization. This data directly impacts penalty weighting in stockout cost evaluations, ensuring alignment with operational priorities.

Together, these assumptions and data preparation steps provide the foundational inputs essential for the model to generate actionable optimization outcomes. Each data element plays a crucial role in capturing operational intricacies and cost dynamics inherent to Armada's supply chain network. For detailed descriptions of these input requirements, Appendix B provides a comprehensive breakdown of format and utilization within the model.

3.4 DC-Level Stockout Probability Estimation

To extend Armada's hub-level stockout prediction model and enable proactive mitigation strategies at the DC level, we developed a structured estimation framework (see Figure 5). This framework integrates demand allocation, inventory visibility, and probabilistic risk assessment across the supply chain network. By estimating DC-level stockout probabilities using upstream hub forecasts and local factors, the framework provides the orchestration recovery engine with the granular risk inputs it needs to evaluate recovery strategies such as expediting, transfers, or event delays.

Figure 5: Stockout Probability Estimation Framework

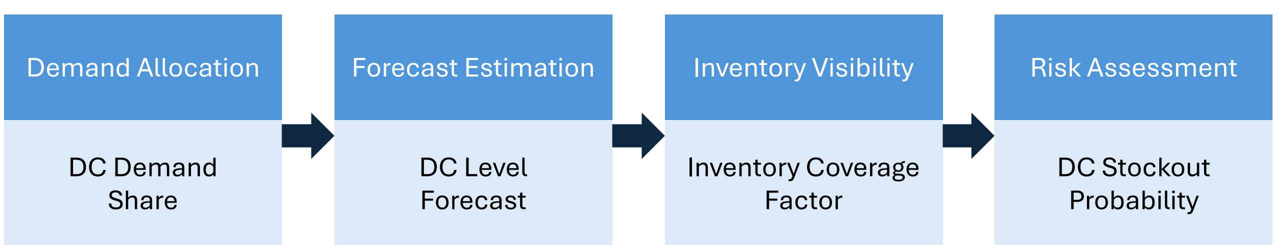

This approach bridges the gap between Armada's upstream forecasts and localized supply chain realities. By translating hub-level insights into DC-level risk estimates, the framework empowers the orchestration recovery engine to make smarter, faster decisions in the face of potential disruptions. To generate DC-level stockout risk estimates from hub-level forecasts, we approximated each DC's expected demand using historical ordering patterns to infer its share of hub-level forecasts. We then evaluated whether each DC had sufficient coverage to meet that demand by comparing available and incoming inventory to projected needs.

The result is a probabilistic estimate of stockout risk at the DC level—derived from upstream signals but weighted by local demand exposure and inventory sufficiency. This estimate enables early, granular visibility into where interventions such as transfers or expedite actions are most needed, ultimately improving the speed and precision of Armada's response strategies.

3.5 Modeling and Optimizing Recovery Actions

Our MILP model is designed to create an optimal transportation plan for a multi-facility supply chain. It integrates multiple cost components and operational constraints to balance competing priorities, minimize costs, and enhance overall network efficiency. Please reference the full MILP optimization model formulation in Appendix A.

3.5.1 MILP Objective Function

The objective function minimizes total cost, which includes transportation costs, early transfer penalty costs, truck underutilization penalty costs and stockout penalty costs. These cost terms are derived from Armada's operational data and planning assumptions, refined in collaboration with supply chain stakeholders. Transportation costs for the model were sourced from historical shipment records and spot rate data, while early transfer penalties, stockout penalties and underutilization penalties are tunable planning parameters used by Armada's logistics team.

Simplified Objective Function:

$$min\,(TC + P_u \times UnderutilizedTrucks + P_e \times EarlyTransfers + P_s \times Stockouts\,)$$

***Where*:**

• TC: Total transportation cost
• P_u: Penalty weight for underutilized trucks
• *UnderutilizedTrucks*: Amount of unused truck capacity
• P_e: Penalty weight for early transfers
• *EarlyTransfers*: Number of early transfer cases
• P_s: Penalty weight for stockouts
• *Stockouts*: Unfulfilled demand (cases)

1. **Transportation Cost:** The model directly includes the shipping transportation cost for new trucks deployed on a route. Notably, existing scheduled routes (representing already-planned movements) are treated as a “free ride” in terms of direct transportation cost. However, even when using these scheduled routes, the model enforces that the destination distribution center (DC) must be replenished within the stockout horizon, ensuring that inventory requirements are met. Furthermore, the model distinguishes between truck types by incorporating temperature-controlled configurations (e.g., Dry, Cooler, Frozen). This adjustment ensures that shipments requiring specific temperature conditions are assigned to appropriate vehicles, thereby safeguarding product integrity and reflecting any cost differentials due to specialized equipment.
2. **Underutilization Penalty:** To promote higher truck utilization, the model penalizes routes where the capacity is not fully used. This penalty incentivizes consolidating shipments to maximize the use of each truck's capacity, reducing the number of partially loaded or empty trucks on the road. This term is calculated by comparing the total capacity deployed (a combination of new and existing trucks) against the actual volume of goods shipped. It is important to note that the calculation of shipped volume leverages key item attributes, specifically, the cases per pallet conversion factor from the item master data. This level of granularity ensures that load calculations account for product-specific dimensions and handling requirements, which is critical when assessing truck utilization.

3. **Early Transfer Penalty:** This penalty is applied to shipments that arrive too early at a destination before they are needed. By imposing this cost, the model discourages premature deliveries that might result in excess inventory or increased handling complexity. It works in tandem with the transportation cost to discourage inefficient routing decisions.
4. **Stockout Penalty:** The stockout penalty reflects the cost of unmet demand when inventory is not available at a downstream facility within the replenishment horizon. This term creates a cost-based tradeoff between allowing a stockout to occur versus triggering a recovery shipment to avoid it. A key feature of the model is its ability to adjust this tradeoff based on business priorities. Specifically, a SKU Classification Matrix and Strategic Customer Facility Matrix are applied to dynamically scale the penalty based on SKU importance and customer criticality. These matrices are derived from flags within the item master and facility master data, such as 'strategic SKU' or 'top-tier customer', allowing the model to assign higher penalties to stockouts that impact high-priority items or customers. By incorporating these differentiated penalty weights, the model supports more strategic replenishment decisions that favor service level preservation where it matters most.

While the model can incorporate additional cost components, such as labor and handling costs based on item attributes like weight, size, or special handling needs, Armada has chosen not to include these in the current scope. If desired in future iterations, a handling cost term could be added to the objective function to assign a cost per case and/or pallet reflecting these operational considerations. For the purposes of this model, shipping and handling costs are excluded.

3.5.2 Cross-Dock (XD) Transshipments Considerations

As explained in the State of Practice chapter, lateral transshipments refer to the movement of inventory between locations operating at the same level of the supply chain. This concept is extended in the model to support **cross-dock (XD) transshipments**, which allow goods to transfer directly between inbound and outbound trucks at intermediate hubs and distribution centers (DCs), bypassing storage and minimizing handling. Armada's optimization model utilizes XD transshipments to consolidate full truckload (FTL) shipments, improving truck utilization, reducing the number of trucks required, and ultimately lowering transportation costs.

Cross-docking has proven to be an effective strategy for boosting logistics network efficiency. For example, a case study in the fast-moving consumer goods (FMCG) retail sector found that shifting from traditional warehousing to a cross-docking distribution approach yielded lower overall supply chain costs (Benrqya, 2019). Such evidence supports the inclusion of cross-docking as an effective tactic to streamline inventory flow and reduce handling requirements in a distribution network.

3.5.2.1 Illustrative Example of Cross-Dock (XD) Benefits

Figure 6 depicts how enabling cross-dock transshipments can reduce total truckloads and improve fleet utilization. In the Cross-Dock Disabled scenario, three separate full truckload (FTL) shipments are required, one of which is highly underutilized, carrying just a single pallet. In contrast, the Cross-Dock Enabled scenario consolidates the underutilized shipment by rerouting it through an intermediate DC, fully utilizing both trucks and reducing the total number of trucks required from three to two. This example highlights the operational efficiency and cost-saving potential of XD transshipments when routing flexibility is available.

Figure 6. Cross-Dock (XD) Transshipment

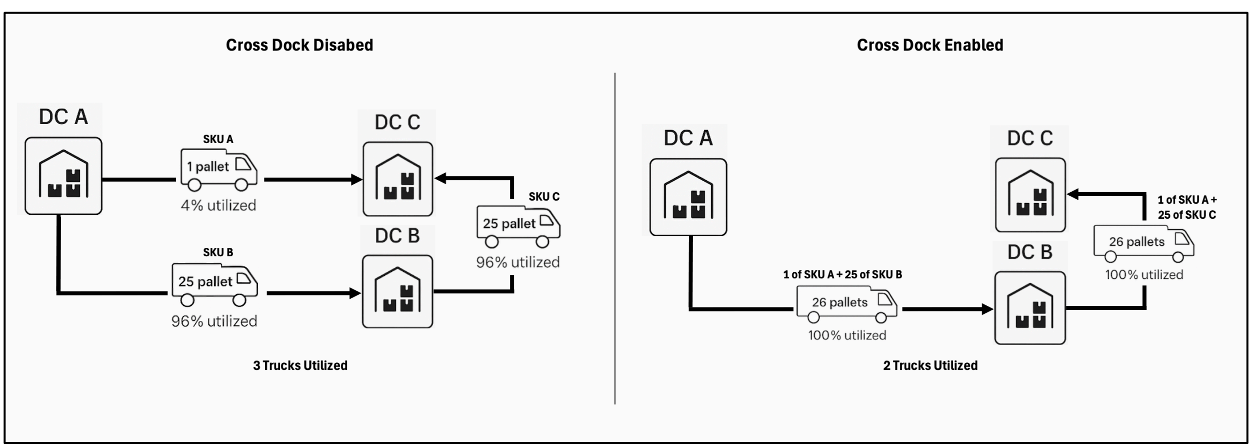

3.5.2.2 Transshipment Cross-Dock (XD) Operational Limitations

Although transshipments can significantly enhance consolidation and utilization, practical constraints within a company's operations could limit their universal applicability. Specifically, the inability to seamlessly execute cross-dock transfers arises from various systemic challenges, including inconsistencies in master data across warehousing systems, complexities in Purchase Order (PO) coordination and scheduling, incremental handling costs, and the organizational challenges associated with partnerships across business units. These considerations were highlighted in discussions with our sponsor company (sponsor company representative, personal communication, February 2025). Recognizing these complexities, the model incorporates a configurable toggle to enable or disable cross-dock transshipments, thus providing flexibility to match operational realities and business-specific constraints.

3.5.3 Interactions Among Cost Levers

The cost levers described in section 3.5.1 are not standalone; they counterbalance one another in the objective function:

1. **Incorporating Early Transfer Penalties:** The early transfer penalty adds another layer of control by preventing the premature delivery of goods. This is critical when aligning shipments with demand windows at the destination DC. The cost incurred by early transfers is weighed against the benefits of consolidated, full-capacity shipments, further tuning the timing of transfers to align with actual demand requirements.
2. **Balancing Transportation, Early Transfer, and Underutilization Costs:** While transportation cost incentivizes the model to minimize the number of new trucks activated, it must be balanced against early transfer and underutilization penalties. The model prioritizes using existing scheduled routes, which are treated as cost-neutral from a transportation standpoint. However, these routes still carry the risk of early arrivals, potentially triggering early transfer penalties if the inventory is delivered ahead of its required window. Simultaneously, the underutilization penalty encourages efficient load planning by penalizing trucks that are not filled to capacity. Together, these cost components create a trade-off: additional trucks

can be deployed when necessary to meet service levels but doing so inefficiently, either by sending trucks too early or by dispatching them with unused capacity, is financially discouraged to promote operational efficiency.

3. **Transshipment Cross-Docks as a Consolidation Strategy:** The model optionally allows transshipments, which can enhance consolidation opportunities. By transferring goods between facilities, transshipments can improve truck utilization. Although these transfers may incur their own early transfer penalty, they are instrumental in reducing the number of empty trucks by creating more efficient load combinations. This trade-off ensures that transshipments are only used when they contribute to an overall reduction in cost.

3.5.4 Operational Constraints

The model is also governed by several critical constraints:

1. **Truck Capacity Constraints:** Every route is constrained by a fixed truck capacity, ensuring that shipments do not exceed the physical limits of the transportation vehicles. Additionally, these constraints are made temperature-aware, meaning that goods requiring specific storage conditions are matched with trucks equipped for that purpose.
2. **Inventory Balance and Replenishment Constraints (Linking Constraint):** The inventory at each facility is tracked over time. The model ensures that initial inventories are correctly set and that the flow of goods between facilities (through direct shipping or transshipments) is balanced by proper inventory updates. The replenishment of the destination DC is explicitly timed to avoid stockouts within a specified recovery horizon.
3. **Existing Scheduled Routes as 'Free Rides':** When a route is part of an already scheduled operation, the transportation cost component is effectively zero. However, these routes still contribute to overall capacity utilization and inventory balance. This feature acknowledges the operational reality that while certain movements are cost-free, they still affect network performance by consuming capacity and influencing timing.

By incorporating shipping transportation costs, early transfer penalties, underutilization penalties, and stockout penalties as interacting levers, the model effectively creates a trade-off scenario. The optimizer must weigh the cost of deploying additional trucks against the benefits of fully utilizing available capacity and timing deliveries to avoid early transfers. Importantly, the stockout penalty introduces a cost-based trade-off between allowing a stockout to occur versus triggering a recovery shipment to prevent it. Constraints such as truck capacity and inventory balance, together with the strategic use of transshipments and scheduled routes, further refine the decision-making process to produce an optimal, cost-effective shipment plan. This integrated approach is particularly useful in complex supply chain networks where multiple objectives and constraints must be balanced concurrently.

3.6 Customizability & Tuning: Bespoke Optimization Levers

The dashboard interface of the optimization model provides decision-makers with a structured and user-friendly mechanism to precisely adjust and calibrate the system through a comprehensive set of configurable input parameters (see section 4.2 and Appendix C for more details on the graphical user interface). These parameters serve as levers, enabling users to align the optimization outcomes with specific strategic priorities and operational realities. Key settings within the dashboard are organized into clear categories for ease of configuration:

1. **Operational Settings:** Users define fundamental constraints of the supply chain network, including the **Time Periods (T),** which establishes the planning horizon, and **Truck Capacity**, specifying the maximum allowable load per transportation unit. Together, these settings frame the scope and scale of optimization.
2. **Penalty Factors:** To balance competing priorities, users can adjust penalty parameters, such as the **Stockout Penalty**, which quantifies the cost associated with unmet demand, the **Underutilization Penalty**, designed to promote efficient use of transportation resources, and the **Early Transfer Penalty**, which discourages premature shipments that could disrupt inventory balance or operational efficiency.

3. **Transshipment Cross-Dock (XD) Toggle:** Decision-makers have the flexibility to activate or deactivate the **Allow XD Transshipments** feature, enabling consideration of intermediate shipment points to enhance consolidation opportunities, optimize truck utilization, and minimize empty mileage.
4. **Priority Weighting:** Recognizing the strategic variability among products and customers, the dashboard offers sophisticated priority weighting matrices. The **SKU Classification Matrix - Priority Weights** (Promotional, Seasonal, & A-B-C SKU Classifications) allows users to assign differentiated levels of importance to product categories based on business impact. Similarly, the **Strategic Customers - Facility Priority Weights Matrix** (High, Medium, & Low) ensures facility-level priorities accurately reflect customer importance and strategic business relationships.

By offering a robust yet intuitive interface, these dashboard configurations empower users to closely tailor the optimization model to meet diverse business needs, ensuring that outcomes align with operational goals, strategic priorities, and real-world constraints. Additionally, the dashboard enables planners to run multiple scenarios and versions, providing insights to support their decision-making process.

4 RESULTS AND DISCUSSION

As outlined in the Methodology chapter, the model's objective function optimizes supply chain decisions by balancing four distinct cost components:

1. **Transportation Cost:** The choice between scheduling a new dedicated truck versus utilizing existing scheduled truck capacity ('free ride') for a given temperature zone and time period.
2. **Truck Underutilization Cost:** Penalties for any unused capacity, applicable to both new and existing scheduled trucks.
3. **Early Transfer Penalty:** Costs incurred when products are moved earlier than necessary (>1 periods ahead of their demand).

4. **Stockout Cost:** Penalties applied if the model either cannot or decides not to replenish a SKU facing an imminent stockout.

Most of these cost elements function as adjustable parameters ('levers' and 'knobs') within the decision-support dashboard, allowing planners to dynamically explore scenarios aligned with their strategic priorities. The model then weighs these cost levers to determine whether a recovery strategy is cost-effective enough to mitigate stockout events or if delaying action or accepting the stockout is the more economical outcome.

To evaluate the behavior of the optimization model under different operational priorities, we conducted a structured scenario analysis. Each scenario isolates and adjusts specific cost levers—such as stockout penalties, truck underutilization, and early transfer costs—to observe how the model responds. The following sections compare the model's recommended recovery strategies across varying configurations, highlighting the trade-offs between service level performance, cost efficiency, and truck utilization.

4.1 Scenario Analysis

In this section, we first establish the conditions of a baseline scenario, which will serve as a reference point for comparing the outcomes of activating various model levers that influence the objective function. The table below summarizes the scenarios analyzed in this section, along with the corresponding input parameters that drive the described outcomes. A summary of the scenario outcomes is provided in Section 4.1.7.

Scenario	Early Transfer Penalty Weight (*Pe*)	Underutilization Penalty Weight (*Pu*)	Stockout Penalty Weight (*Ps*)	Cross-Dock (XD) Enabled *(Yes/No)*
Baseline (4.1.1) & Transportation Cost Dominant (4.1.2)	*$1*	*$0.50*	*$15*	*No*
Truck Utilization Dominant (4.1.3)	*$1*	***$5***	*$15*	*No*
Early Transfer Penalty Dominant (4.1.4)	***$20***	*$0.50*	*$15*	*No*
Stockout Cost Dominant (4.1.5)	*$1*	*$0.50*	***$9,999*** *(Prevents stockouts unless unavoidable)*	*No*
Cross-Dock (XD) Transshipment (4.1.6)	*$1*	*$0.50*	*$15*	***Yes***

4.1.1 Baseline Scenario

In the baseline scenario, the model evaluates on-hand inventory at the locations listed below as eligible source points for transfers and replenishments to prevent stockouts across all DCs included in this results section. Projected stockout events are also outlined below, along with a scheduled cooler truck route from **DC A** to **DC B**, representing a potential “free ride” for cooler pallets in the model. This route already has 2 of its 26 pallet spaces allocated, leaving 24 available for potential transfers. For simplicity, our initial scenarios only consider lateral transshipments between DCs, though in practice, inventory could be sourced from any eligible warehouse, including hubs or nearby DCs.

On Hand Inventory Position in Network Prior to Recovery Events (If not listed, assume 0 inventory for DC/SKU combination)

DC	SKU	On Hand Position
DC A	**SKU A**	10 Pallets
DC A	**SKU B**	10 Pallets

Stockout Demand Signals

Time Period (t) of Projected Stockout Event	DC	SKU + Qty	Temp	Planning Status
1	DC B	10 Pallets of SKU A	Cooler	*Shipment Unplanned*
4	DC B	10 Pallets of SKU B	Cooler	*Shipment Unplanned*

Existing Scheduled Transportation ('Free Ride')

Time Period (t) of 'Free Ride' Existing Scheduled Route	Origin DC	Dest DC	Temp	Total Truck Capacity	Claimed (Pre-Planned) Capacity	Residual Capacity
1	DC A	DC B	Cooler	26 Pallets	*2 Pallets*	24 Pallets

In the baseline scenario, our sponsor company currently captures projected stockout risk at the hub level, then conducts follow-up analysis to identify which downstream DCs are most impacted. Replenishment or transfer decisions are made manually on a case-by-case basis to mitigate the risk. In the following scenarios, we demonstrate how our model generates automated, optimized recommendations, and how the various components of the objective function, along with tunable input parameters, influence the nature of those recommendations.

4.1.2 Transportation Cost Dominant Scenario

In a Transportation Cost Dominant scenario, the model prioritizes reducing truck transportation costs over any other cost levers in the objective function, such as early penalty costs. Consequently, it fully utilizes the scheduled truck in period 1 by shipping pallets of both SKU A and SKU B simultaneously. This allows the model to take advantage of an existing scheduled route which is a cost-neutral option, largely reducing transportation costs. In this scenario, the model will penalize the planner with an early transfer penalty, however an overall cost reduction is realized (see section 4.1.7 for full scenario analysis summary).

Transportation Cost Dominant Scenario

Time Period (t) of Replenishment	Origin DC	Dest DC	Trucks	SKU + Qty	Truck Utilization	Temp
1 (*Existing Route*)	DC A	DC B	1 Truck	10 Pallets of SKU A 10 Pallets of SKU B *2 Pre-Planned Pallets*	85% (22 / 26)	Cooler

Dominant Term in Objective Function:

$$min\ (\ TC + \ldots)$$

***Where*:**

• TC: Total transportation cost

4.1.3 Truck Utilization Penalty Dominant Scenario

In this scenario, the model prioritizes reducing truck underutilization penalties over early transfer penalties. Consequently, it fully utilizes the scheduled truck in period 1 by shipping pallets of both SKU A and SKU B simultaneously, same as the prior transportation cost dominant scenario (4.1.2). While leveraging scheduled routes may appear cost-efficient, particularly when adding shipments to an already-planned truck as a 'free ride', the model must still account for underutilization penalties. When the corresponding parameter is set to heavily discourage empty capacity, the cost of unused space can outweigh the benefits of minimizing transportation or early transfer costs, shifting the model toward fuller trailer utilization.

Truck Utilization Penalty Dominant Scenario

Time Period (t) of Replenishment	Origin DC	Dest DC	Trucks	SKU + Qty	Truck Utilization	Temp
1 (*Existing Route*)	DC A	DC B	1 Truck	10 Pallets of SKU A 10 Pallets of SKU B *2 Pre-Planned Pallets*	85% (22 / 26)	Cooler

Dominant Term in Objective Function:

$$min\,(\ldots\ +\ P_u\ \times\ UnderutilizedTrucks\ +\ \ldots\,)$$

***Where*:**

• P_u: Penalty weight for underutilized trucks
• *UnderutilizedTrucks*: Amount of unused truck capacity

This choice maximizes truck utilization by consolidating multiple SKUs onto a single scheduled route, but it also introduces potential risks and added costs associated with the premature movement of SKU B, which may arrive before it is needed and trigger early transfer penalties or excess handling at the destination. This trade-off is quantified in the model results by applying an early transfer penalty to the resulting objective function output.

4.1.4 Early Transfer Penalty Dominant Scenario

In this scenario, the model emphasizes minimizing early transfer penalties. Despite available truck space in period 1, the model opts to delay SKU B shipment until closer to its predicted stockout event, resulting in two separate shipments, each under 50% utilization.

Early Transfer Penalty Dominant Scenario

Time Period (t) of Replenishment	Origin DC	Dest DC	Trucks	SKU + Qty	Truck Utilization	Temp
1 *(Existing Route)*	DC A	DC B	1 Truck	10 Pallets of SKU A *2 Pre-Planned Pallets*	46% (12 / 26)	Cooler
4	DC A	DC B	1 Truck	10 Pallets of SKU B	38% (10 / 26)	Cooler

Dominant Term in Objective Function:

$$min\,(\ldots\ +\ P_e \times EarlyTransfers\ +\ \ldots\ )$$

***Where*:**

• P_e: Penalty weight for early transfers
• *EarlyTransfers*: Number of early transfer cases

This strategy reduces the risk and potential costs associated with premature product movements, aligning replenishment closely with stronger stockout signals.

4.1.5 Stockout Cost Dominant Scenario

When stockout penalties are set lower relative to the combined penalties of truck underutilization, early transfers, and transportation, the model might suggest stockouts as a cost-effective option. To prevent stockouts from being recommended except under unavoidable circumstances (insufficient inventory), planners could in theory assign a high stockout penalty per case (e.g., $99,999), compelling the model to exhaust all other feasible options before accepting a stockout event (see section 4.1.7 for full scenario analysis summary).

Dominant Term in Objective Function:

$$min\,(\ldots\ +\ P_s\ \times\ Stockouts\,)$$

***Where*:**
• P_s: Penalty weight for stockouts
• *Stockouts*: Unfulfilled demand (*cases*)

Based on discussions with Armada's logistics and planning teams, the average estimated cost of a stockout is approximately $15-18 per case, depending on the product category and customer impact. As a starting point, our model applies a baseline stockout penalty of $15 per case, which is a configurable parameter in the front-end user dashboard. This base value is dynamically adjusted using two additional configurable multipliers: the SKU Classification Matrix and the Strategic Customer Facility Matrix. These allow the stockout cost to scale upward when high-priority SKUs or strategic customer facilities are impacted, based on designations in the item master and facility master data. This structure ensures that the optimization engine aligns replenishment decisions with business-critical priorities and real-world service level expectations.

4.1.5.1 Priority-Based Stockout Penalty Weights

In the model results, we observed that the stockout penalty dynamically increased for SKUs and facilities based on their assigned priority. This behavior is driven by two layers of weighting:

1. **SKU Classification Matrix**: Items are classified as P-S-A-B-C (i.e., Promotional, Seasonal, or A-B-C velocity), and the higher priority classifications (Promotional, seasonal, and A velocity) are considered more expensive to stockout.
2. **Strategic Customer / Facility Priority:** where shipments to high-priority customers or facilities are similarly penalized more severely in the event of a stockout.

These weights are defined by the planner during model setup and are used to multiply the base stockout penalty. The model's output reflected this prioritization by consistently favoring fulfillment of high-priority items and customers when trade-offs had to be made.

4.1.5.1.1 Observed Weight Multipliers

SKU Classification Stockout Weight Multipliers

SKU Priority	P (Highest)	S	A	B	C (Lowest)
Weight	1.20	1.15	1.10	1.05	1.00

Facility/Customer Priority Stockout Weight Multipliers

Facility Priority	High	Medium	Low
Weight	1.20	1.10	1.00

4.1.5.1.2 Resulting Impact

If the base stockout penalty is set to 15 per case, and an item is Priority A (1.10), going to a Medium-priority facility (1.10), the effective stockout penalty applied by the model becomes:

$$\mathbf{\$15 \times 1.10 \times 1.10 = \underline{\$18.15}}$$

As a result, the model output clearly reflected a strong bias against stockouts for:

1. Promotional, seasonal and high velocity SKUs
2. High-priority facilities or strategic customer locations

This weighted penalty system allowed the model to automatically steer fulfillment toward the most critical areas of the network, ensuring higher service levels where it matters most, and only accepting stockouts where recovery costs were prohibitive or infeasible.

4.1.6 Cross-Dock (XD) Transshipment Scenario

An additional feature within the model is the "Allow XD Transshipments" toggle. When enabled, the model can route products through an intermediate facility (i.e., cross-docking), creating the opportunity for lower-cost fulfillment options by chaining together transportation legs across facilities. This option is especially powerful when direct shipment options are costly or unavailable, allowing the model to optimize network flow based on total cost, inventory availability, and truck utilization. However, when cross-docking is disabled, the model is restricted to direct shipments only. As a result, it often has fewer low-cost options, and in many cases, may determine that stockouts are more cost-effective than fulfilling demand at a cost that exceeds the stockout penalty.

4.1.6.1 Cross-Dock Baseline Scenario

The previous baseline scenario from 4.1.1 has been extended into a cross-dock baseline scenario to include a projected stockout event at **DC C**. In this updated scenario, the model continues to evaluate the on-hand inventory at the locations listed below as eligible source points for transfers and replenishments to mitigate stockouts. Projected stockout events are also shown below, along with the same scheduled cooler truck route from **DC A** to **DC B**, representing a potential "free ride" for cooler pallets in the model. This route currently has 2 of its 26 pallet spaces allocated, leaving 24 available for potential transfers.

On Hand Inventory Position in Network Prior to Recovery Events (If not listed, assume 0 inventory for DC/SKU combination)

DC	SKU	On Hand Position
DC A	**SKU A**	10 Pallets
DC A	**SKU B**	10 Pallets
DC A	**SKU C**	12 Pallets
DC B	**SKU D**	12 Pallets

Stockout Demand Signals

Time Period (t) of Projected Stockout Event	DC	SKU + Qty	Temp	Planning Status
1	DC B	10 Pallets of SKU A	Cooler	*Shipment Unplanned*
4	DC B	10 Pallets of SKU B	Cooler	*Shipment Unplanned*
2	DC C	12 Pallets of SKU C	Cooler	*Shipment Unplanned*
2	DC C	12 Pallets of SKU D	Cooler	*Shipment Unplanned*

Existing Scheduled Transportation ('Free Ride')

Time Period (t) of 'Free Ride' Existing Scheduled Route	Origin DC	Dest DC	Temp	Total Truck Capacity	Claimed (Pre-Planned) Capacity	Residual Capacity
1	DC A	DC B	Cooler	26 Pallets	*2 Pallets*	24 Pallets

Route Feasibility Map

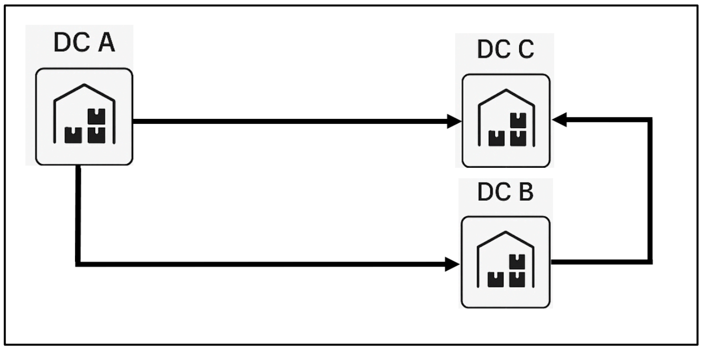

4.1.6.2 With Cross-Docking Enabled

The model identifies a cost-effective transshipment path from DC A to DC C via DC B by leveraging existing transportation capacity and avoiding a high-cost, underutilized direct shipment. This approach is especially valuable when direct shipment options are limited or expensive, enabling the model to optimize network flow based on total cost, inventory availability, and truck utilization. Based on initial model runs, enabling cross-dock transshipments at scale may yield transportation cost savings of up to 15%, depending on scenario configuration and baseline truck utilization, primarily by reducing the need for additional direct shipments and increasing efficiency across multi-leg routes.

Cost-Effective XD Transshipment Using Existing Capacity

Time Period (t) of Replenishment	Origin DC	Dest DC	Trucks	SKU + Qty	Truck Utilization	Temp
1 *(Existing Route)*	DC A	DC B	1 Truck	10 Pallets of SKU A 12 Pallets of SKU C *2 Pre-Planned Pallets*	92% (24 / 26)	Cooler
2	DC B	DC C	1 Truck	12 Pallets of SKU C 12 Pallets of SKU D	92% (24 / 26)	Cooler
4	DC A	DC B	1 Truck	10 Pallets of SKU B	38% (10 / 26)	Cooler

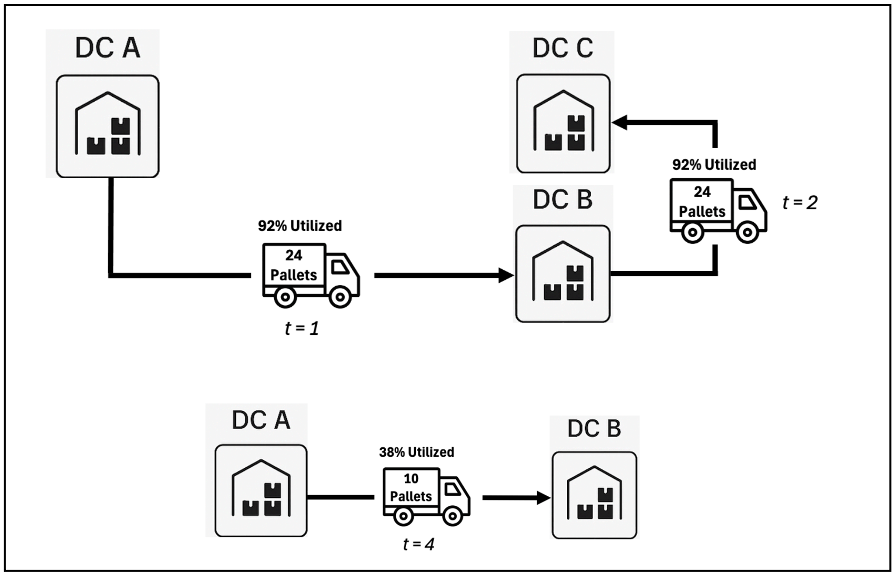

4.1.6.3 With Cross-Docking Disabled

With cross-docking disabled, the model restricts the ability to flow inventory through a DC via cross-dock transshipments, which will either nudge the model towards either shorting the demand (thus creating a stockout event) or offering a more expensive recovery strategy. This would depend on the Stockout Penalty parameters set by the planners. If we were to assume the above scenario does not recommend a stockout, then the recommended shipment plan would be as follows:

XD Transshipment Disabled

Time Period (t) of Replenishment	Origin DC	Dest DC	Trucks	SKU + Qty	Truck Utilization	Temp
1 *(Existing Route)*	DC A	DC B	1 Truck	10 Pallets of SKU A *2 Pre-Planned Pallets*	46% (12 / 26)	Cooler
2	DC A	DC C	1 Truck	12 Pallets of SKU C	46% (12 / 26)	Cooler
2	DC B	DC C	1 Truck	12 Pallets of SKU D	46% (12 / 26)	Cooler
4	DC A	DC B	1 Truck	10 Pallets of SKU B	38% (10 / 26)	Cooler

Note: t=4 shipment could consolidate on t=1 shipment, but would accrue Early Transfer Penalty

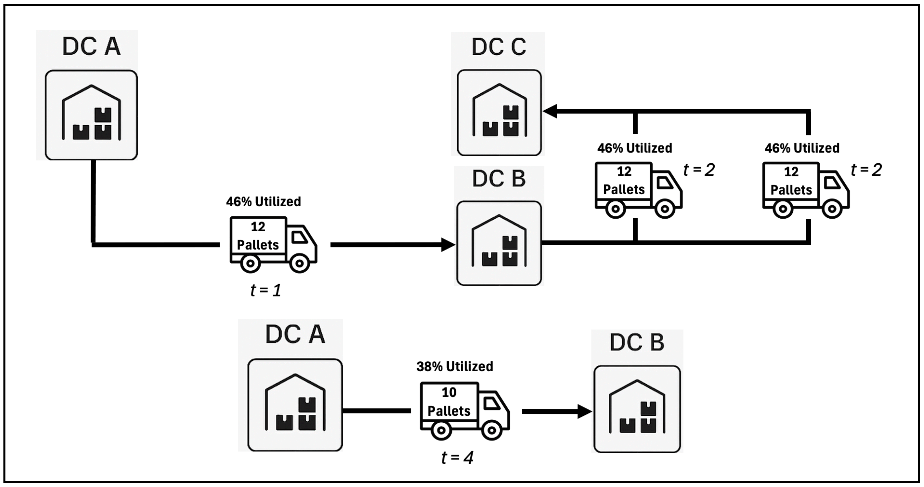

The inefficiencies observed in the above example highlight the flexibility and value of enabling cross-docking. When it is enabled, the model can dynamically assemble fulfillment strategies across the broader network. When disabled, options become more limited, often resulting in higher-cost recovery strategies or, in some cases, stockouts, due to the prohibitive cost of executing a recovery event, even when inventory exists elsewhere in the network.

4.1.7 Scenario Analysis Summary

The scenario analyses illustrate how different terms in the model's objective function interact and demonstrate how the model dynamically adapts to shifting cost priorities by adjusting replenishment strategies based on the relative weightings of transportation cost, truck underutilization, early transfer penalties, and stockout risk.

Scenario	Early Inventory Transfers	Net Truck Utilization	Stockouts	Transportation Cost Reduction
Baseline (4.1.1) & Transportation Cost Dominant (4.1.2)	↑ *(Accepted early)*	↑ *(85%)*	↓	↓ *(Reduced by free ride)*
Truck Utilization Dominant (4.1.3)	↑	↑ *(85%)*	↓	↓↓
Early Transfer Penalty Dominant (4.1.4)	↓ *(Delayed)*	↓ *(42%)*	↓	↑↑ *(Higher due to split loads)*
Stockout Cost Dominant (4.1.5)	↑ *(Forced if inventory exists)*	↓↓ *(Low use to avoid stockout)*	↓↓ *(Always avoided, unless infeasible)*	↑ *(May be high to avoid stockouts)*
Cross-Dock (XD) Transshipment (4.1.6)	↑	↑ *(92%)*	↓	↓↓↓ *(Best savings with XD)*

When transportation cost or truck utilization is prioritized, the model consolidates shipments to maximize use of existing routes, even if it means shipping products earlier than needed. In contrast, when early transfer penalties are emphasized, the model delays shipments to avoid premature arrivals, often at the expense of truck efficiency. Scenarios with lower stockout penalties may even result in accepted shortages when recovery options are deemed too costly. Enabling cross-dock transshipments significantly expands the solution space, unlocking cost-effective routing strategies that improve truck utilization and reduce reliance on direct shipments. Taken together, these scenarios highlight the model's flexibility in evaluating trade-offs and guiding planners toward recovery strategies aligned with cost, service, and operational goals.

4.2 Dashboard Model Graphical User Interface (GUI)

The model's graphical user interface (GUI) dashboard enables end-users to input a wide range of planning parameters, such as stockout penalty, truck underutilization penalty, and early transfer penalty, and rerun the optimization model as needed. This interactive functionality allows planners to explore recommended supply chain strategies across different cost scenarios, supporting data-driven decisions that align with evolving business priorities and operational realities.

To support interpretation and usability, additional charts and visualizations have been incorporated into the dashboard. These graphics, in Appendix C, provide insights into shipment timing, truck utilization, inventory flows, and projected stockout risks. By presenting this information in a visual format, the dashboard enables faster analysis, clearer communication, and more informed decision-making.

Additional capabilities have also been integrated to enhance the dashboard's visual and functional usability. Recommended truckloads are displayed in a map-based format using Folium, allowing planners to quickly visualize freight movements across the network. A Shipment Planning Matrix has been introduced to provide detailed, time-phased freight recommendations by SKU, quantity, origin, and destination. To further support operational execution and analysis outside the platform, users can download the complete set of planning recommendations directly in Excel format via built-in export functionality.

4.3 Model Limitations & Areas for Improvement

While the optimization model provides a robust mechanism for improving transportation efficiency and stockout management, several limitations have emerged, indicating opportunities for future enhancement or exploration:

1. **Stockout Prediction Granularity:** The current model architecture requires DC-level stockout predictions as an input. However, Armada's existing predictive capabilities generate stockout forecasts only at the SKU and Hub levels. This limitation necessitates an intermediate step to translate these Hub-level predictions into actionable DC-level insights. To improve model precision and operational applicability, it is recommended that additional efforts be directed toward developing advanced, machine-learning-driven stockout prediction models specifically at the DC level, facilitating direct and precise inputs for the optimization model.
2. **Lead Time Simplification:** In the current research, the model employs standardized transportation lead times between Hubs and DCs, regardless of specific origin-destination combinations. Transportation lead times often vary significantly due to geographical distance, carrier reliability, and traffic patterns. Future iterations of the model could incorporate differentiated lead times based on historical transit data and route-specific operational insights, thereby refining replenishment timing and inventory positioning.
3. **Existing Truck Cubing & Load Optimization:** An area for potential enhancement involves optimizing the utilization, or "cubing", of already scheduled replenishment trucks. Currently, the model does not account for the ability to physically substitute pallets containing higher-priority SKUs, particularly those at risk of imminent stockouts, into trucks that have already been planned for Hub-to-DC shipments. This strategic "SKU swapping" could improve responsiveness to urgent stockouts without incurring additional transportation costs. Implementing this approach would require tight operational coordination at the Hubs to ensure pallet substitutions occur before items are picked, staged, or loaded.

4. **Penalty Cost Determination and SKU Classification:** Penalty costs in the model are currently weighted based on a predefined SKU-Classification Matrix (A-B-C-P-S). While this approach effectively prioritizes SKU replenishment, there remains considerable scope for deeper analysis into determining optimal SKU classifications. Future studies should investigate advanced methodologies for dynamically classifying SKUs based on changing market conditions, promotional cycles, or strategic goals. Additionally, quantifying the true business cost associated with stockouts at the SKU level, beyond mere product prioritization, warrants deeper exploration.
5. **Strategic Customer Prioritization:** While the model currently includes facility-level strategic customer prioritization, additional granularity could further enhance its strategic alignment. Expanding the prioritization logic to encompass additional customer-related shipping attributes, such as downstream customers ("customer's customer") or specific delivery commitments, would allow the model to more accurately reflect strategic business objectives and customer service requirements.
6. **Trailer Temperature Management Constraints:** The current model simplifies temperature-controlled transportation by assuming single-temperature-zone trucks. Real-world logistics, however, frequently utilize multi-temperature-zone trailers, enabling consolidation of dry goods alongside chilled or frozen products. Addressing this complexity within future iterations of the model could unlock additional cost-saving opportunities and logistical efficiencies, aligning optimization results more closely with operational realities.
7. **Model Computational Performance:** The complexity inherent in the optimization framework, particularly the high number of decision variables, introduces computational limitations. Large-scale network optimization scenarios could face challenges in computational performance and scalability. Continued research into algorithmic efficiencies, heuristic approaches, or parallelized computational techniques will be beneficial to address performance bottlenecks and enable faster decision-making.

Addressing these limitations and areas for improvement will refine the robustness, scalability, and strategic value of the optimization model, ensuring its alignment with real-world operational constraints and business objectives.

5 CONCLUSION

This project demonstrates how integrating predictive insights with optimization modeling can proactively mitigate stockouts and strengthen supply chain resilience in complex distribution networks. For management, we recommend adopting the recommendation engine as part of the standard planning cadence to evaluate trade-offs between cost, urgency, and service level across a range of recovery strategies. The dashboard interface enables planners to fine-tune model behavior and explore scenario-based decisions that align with business priorities. Scenario analysis capabilities further enable teams to stress-test planning assumptions, visualize cost-performance trade-offs, and identify the most cost-effective fulfillment strategies under varying conditions. We also encourage ongoing investment in upstream data quality and the continued refinement of stockout predictions at the DC level to increase the precision of recovery actions.

Assuming on-hand inventory is available within the replenishment network and reachable within acceptable lead times, the model has demonstrated potential to reduce total stockouts to 0%. However, when inventory is not positioned accessibly or transfer costs are too high, recovery actions may become cost prohibitive. In such cases, the model helps planners quickly determine whether to trigger a recovery event or accept a stockout. Compared to manual planning, the model offers a faster, more consistent, and analytically rigorous way to navigate recovery decisions, allowing planners to modulate decision levers, compare trade-offs, and choose informed actions that align with operational and financial goals.

Looking ahead, future development should focus on incorporating dynamic transportation lead times informed by route-specific data and enabling in-transit adjustments, such as reprioritizing pallets on scheduled Hub-to-DC shipments to better support stockout mitigation and adapt to evolving network conditions. Additionally, we recommend integrating inventory optimization methods, including safety stock management and inventory balancing, to reduce reliance on reactive transfers and enhance overall network efficiency. Addressing these areas will improve operational flexibility and support scalable deployment across Armada’s end-to-end supply chain.

While this project was developed in collaboration with Armada, the underlying framework is applicable to other organizations managing multi-tier distribution networks. Companies seeking to improve recovery planning, reduce stockouts, and enhance supply chain responsiveness can adapt this approach by integrating predictive insights with cost-aware, scenario-based optimization models.

REFERENCES

Adoolffo, G. F. R. D. S., Baierle, I. C., Gomes, L. D. C., Correa, R. G. D. F., & Peres, F. A. P. (2023). A comprehensive Roadmap for Connecting Industry 4.0 Technologies to the Basic Model of Collaborative Planning, Forecasting, and Replenishment (CPFR). *MDPI*. https://www.mdpi.com/2076-3387/14/6/108

Axsäter, S (2015). Inventory Control. *International Series in Operations Research & Management Science.* https://link.springer.com/book/10.1007/978-3-319-15729-0

Benrqya, Y. (2019). Costs and benefits of using cross-docking in the retail supply chain: A case study of an FMCG company. *International Journal of Retail & Distribution Management*, 47(4), 412–432. https://doi.org/10.1108/IJRDM-07-2018-0119

Bertrand, L. P., & Bookbinder, J. H. (1998). Stock redistribution in two-echelon logistics systems. *Journal of the Operational Research Society.* https://link.springer.com/article/10.1057/palgrave.jors.2600613

Brunaud, B., Laínez-Aguirre, J. M., Pinto, J. M., & Grossmann, I. E. (2018). Inventory policies and safety stock optimization for supply chain planning. *AIChE Journal.* https://aiche.onlinelibrary.wiley.com/doi/full/10.1002/aic.16421

Caplice, C., & Ponce, E. (2024). MITx MicroMasters program in supply chain management: Key concepts. *MIT.* https://courses.edx.org/asset-v1:MITx+CTL.SC1x_2+1T2016+type@asset+block/SC1x_KeyConceptDocument_v5_1_Complete.pdf

Cao, Z., Yang, Z., Liu, F., & Ye, Z. (2018). Buyback and stockout penalty newsvendor model based on reference point prospect theory. *Springer Science+Business Media, LLC.* https://link-springer-com.libproxy.mit.edu/article/10.1007/s10586-018-2451-1

Driessen, M. (2023). Unlock the power of multi-echelon inventory optimization. *Journal of Business Forecasting.* https://eyeonplanning.com/wp-content/uploads/2024/02/Multi-Echelon-Inventory-Optimization.pdf

Eppen, G. D. (1979). Effects of centralization on expected costs in a multi-location newsboy problem. *Management Science,* 25(5), 498-501. https://www.jstor.org/stable/2630280

Halbe, H., & Je, J. (2022). Inventory Rebalancing through Lateral Transshipments. Master of Applied Science capstone, *Massachusetts Institute of Technology.* https://dspace.mit.edu/handle/1721.1/142952

Hill, C. A., Zhang, G. P., & Miller, K. E. (2018). Collaborative Planning, Forecasting, and Replenishment & Firm Performance: An Empirical Evaluation. *International Journal of Production Economics.* https://doi.org/10.1016/j.ijpe.2017.11.012

Lee, Y. H., Jung, J. W., & Jeon, Y. S. (2006). An Effective Lateral Transshipment Policy to Improve Service Level in the Supply Chain. *International Journal of Production Economics.* https://doi.org/10.1016/j.ijpe.2006.05.007

Moin, N., & Salhi, S. (2007). Inventory Routing Problems: A Logistical Overview. *The Journal of the Operational Research Society, 58(9), 1185-1194.*

Nigam, A. (2016). Product promotion effectiveness: Root causes of stockouts [Thesis, Massachusetts Institute of Technology]. *MIT DSpace.* https://dspace.mit.edu/handle/1721.1/107513

Oeser, G. (2015). Risk-Pooling Essentials. *Springer International Publishing.* https://link.springer.com/book/10.1007/978-3-319-14157-2

Panahifar, F., Heavey, C., Byrne, P. J., & Fazlollahtabar, H. (2015). A framework for Collaborative Planning, Forecasting and Replenishment (CPFR). *Journal of Enterprise Information Management.* https://doi.org/10.1108/JEIM-09-2014-0092

Sharma, H. K., & Bojorquez Aispuro, A. (2020). Evaluating inventory risk pooling strategy for multi-echelon distribution network (master's thesis). *Massachusetts Institute of Technology.* https://dspace.mit.edu/handle/1721.1/126496

Silver, E. A., Pyke, D. F., & Thomas, D. J. (2017). Inventory Production Management in Supply Chains. Boca Raton: *CRC Press.*

Soysal, M., Bloemhof-Ruwaard, J., Haijema, R., & van der Vorst, J. (2015). Modeling an Inventory Routing Problem for perishable products with environmental considerations and demand uncertainty. *International Journal of Production Economics,* 164, 118-133.

Zhang, Y., Chai, Y., & Ma, L. (2021). Research on multi-echelon inventory optimization for fresh products in supply chains. *School of Automation, Chongqing University.* https://www.mdpi.com/2071-1050/13/11/6309

APPENDICES

Appendix A

Sets and Indices

- $i, j \in \Gamma$: Set of facilities (hubs and DCs)
- $s \in S$: Set of SKUs
- $t \in T$: Set of discrete time periods
- $m \in M$: Set of temperature modes
- $S^m \subseteq S$: Subset of SKUs that require temperature mode m

Parameters

- D_{js}: Demand for SKU s at facility j over the recovery horizon
- S_{is}: Initial supply of SKU s at facility i
- R_{ij}: Transportation cost from facility i to j
- p_s: Cases per pallet for SKU s
- C: Truck capacity (in pallets)
- P_s: Stockout penalty for SKU s
- β: Underutilization penalty per percent unused capacity
- F: Early transfer penalty per case
- E_{ijtm}: Number of existing trucks available on route $i \to j$ at time t and temperature mode m
- H_{js}: Recovery horizon (latest acceptable fulfillment time)

Decision Variables

- $x_{ijst} \in \mathbb{Z}_+$: Cases of SKU s shipped from i to j at time t
- $y^n_{ijtm} \in \mathbb{Z}_+$: Number of new trucks on route $i \to j$ at t, mode m
- $y^e_{ijtm} \in \mathbb{Z}_+$: Number of existing trucks used on same route
- $u_{ijtm} \in \mathbb{R}_+$: Underutilization on route $i \to j$ at t, mode m
- $e_{jst} \in \mathbb{Z}_+$: Early transfer of SKU s to j at t
- $z_{js} \in \mathbb{R}_+$: Stockout quantity for SKU s at j
- $I_{jst} \in \mathbb{Z}_+$: Inventory of SKU s at j at time t

Objective Function

Minimize total cost:

$$
\begin{aligned}
\min \quad & \sum_{i\in\Gamma}\sum_{j\in\Gamma}\sum_{t\in T}\sum_{m\in M} R_{ij}\cdot y^n_{ijtm} \\
& + \sum_{i\in\Gamma}\sum_{j\in\Gamma}\sum_{t\in T}\sum_{m\in M} \beta\cdot u_{ijtm}\cdot \left(y^n_{ijtm}+y^e_{ijtm}\right)\cdot 100 \\
& + \sum_{j\in\Gamma}\sum_{s\in S}\sum_{t\in T} \delta_{jst}\cdot F\cdot e_{jst} \\
& + \sum_{j\in\Gamma}\sum_{s\in S} P_s\cdot z_{js}
\end{aligned}
$$

where:

$$
\delta_{jst} = \begin{cases} 1 & \text{if } t < H_{js}-1 \\ 0 & \text{otherwise} \end{cases}
$$

Constraints

subject to:

1. **Truck capacity per temperature mode**

$$
\sum_{s\in S^m} \frac{x_{ijst}}{p_s} \le C\cdot\left(y^n_{ijtm}+y^e_{ijtm}\right) \quad \forall i,j\in\Gamma,\ t\in T,\ m\in M
$$

2. **Existing truck limit**

$$
y^e_{ijtm} \le E_{ijtm} \quad \forall i,j\in\Gamma,\ t\in T,\ m\in M
$$

3. **Inventory balance**

$$
I_{jst} = I_{jst-1} + \sum_{i\in\Gamma} x_{ijst-1} - \sum_{k\in\Gamma} x_{jkst-1} \quad \forall j\in\Gamma,\ s\in S,\ t\in T\setminus\{0\}
$$

4. **Initial inventory**

$$
I_{js0} = S_{js} \quad \forall j\in\Gamma,\ s\in S
$$

5. **Shipping cannot exceed inventory**

$$
\sum_{k\in\Gamma} x_{jkst} \le I_{jst} \quad \forall j\in\Gamma,\ s\in S,\ t\in T
$$

6. **Demand satisfaction with stockouts**

$$
I_{js,H_{js}} + z_{js} = D_{js} \quad \forall j\in\Gamma,\ s\in S
$$

7. **Post-horizon inventory stability**

$$
I_{jst} = I_{jst-1} \quad \forall j\in\Gamma,\ s\in S,\ t > H_{js}
$$

8. **No transshipment**

$$
x_{ijst} = 0 \quad \text{if } j \notin \text{final destinations of } s,\ \forall i,j\in\Gamma,\ s\in S,\ t\in T
$$

where:

$$
e_{jst} = \begin{cases} \sum_{i\in\Gamma} x_{ijst} & \text{if } t < H_{js}-1 \\ 0 & \text{otherwise} \end{cases}
$$

$$
u_{ijtm} = \frac{C\cdot\left(y^n_{ijtm}+y^e_{ijtm}\right) - \sum_{s\in S^m}\frac{x_{ijst}}{p_s}}{C}
$$

Appendix B

The Item input file outlines key product characteristics required by the model, including product number, temperature type, classification, and UOM conversion logic (cases per pallet). These attributes enable accurate capacity planning, temperature-specific routing, and SKU-level prioritization during optimization.

Figure B.1: Item import file structure used to define product attributes and UOM conversions

PRODUCT_NUMBER	TEMP	CLASS	CASE_QUANTITY	CASES_PER_PALLET
SKU001	FREEZER	C	44	40
SKU002	COOLER	C	288	160
SKU003	DRY	C	12	30
100292	DRY	C	12	80

The Facilities file specifies key attributes for each facility, including facility type (e.g., Hub or DC), unique identifier, geographic coordinates, and classification tier. These inputs support routing logic, distance-based calculations, and priority weighting in the optimization model.

Figure B.2: Facilities import file structure used to define nodes in the warehousing network

type	dc_id	lat	lon	class
Hub	HUB001	...	...	Low
DC	DC001	...	...	Low
DC	DC002	...	...	Low
DC	DC003	...	...	Low

The Transportation file provides the cost matrix used by the optimization model to evaluate transportation expenses across origin-destination pairs. These values directly influence routing decisions by enabling cost-based comparisons between direct and transshipment options.

Figure B.3: Transportation import file structure defining shipment costs between facility pairs

Origin	Destination	Cost
DC001	DC002	625.56
DC001	DC003	1377.96
DC002	DC003	729.35

The Existing Truck file identifies trucks already scheduled to travel between facilities in specified time periods, along with their temperature configuration. These inputs allow the model to incorporate available capacity as a cost-neutral option during optimization.

Figure B.4: Existing truck import file structure used to define pre-scheduled truck movements

origin	destination	time_period	temp_type
DC001	DC002	3	DRY
DC001	DC003	3	DRY
DC002	DC003	4	DRY

The Demand file lists forecasted demand at risk of stockout, including the total quantity required and the recovery time window. These entries enable the model to prioritize and time replenishments based on urgency and demand coverage.

Figure B.5: Stockout demand import file structure defining recovery requirements by facility and SKU

FacilityId	ProductNumber	TotalDemand	RecoveryHorizon
DC001	SKU001	24	6
DC001	SKU002	24	6
DC001	SKU003	96	3
DC002	SKU002	18	2

The On-Hand Position file provides visibility into current inventory levels at each facility by SKU, specifically for SKUs impacted by potential stockout events. It serves as the starting point for the optimization model to identify eligible source locations for recovery shipments.

Figure B.6: Available on-hand import file structure used to represent inventory positions across the network

FacilityId	ProductNumber	OnHand
DC002	SKU001	300
DC002	SKU003	144
DC003	SKU002	336

Appendix C

Figure C.1 displays the interactive dashboard interface used to configure key model parameters, including penalties for stockouts, early transfers, and underutilization. Users can adjust SKU and customer priority weights, toggle cross-docking options, and rerun the optimization based on updated inputs. The dashboard also includes options to reset parameters to default values and export model outputs directly into Excel for further analysis or operational execution.

Figure C.1: Interactive dashboard interface for adjusting model parameters and priority weights

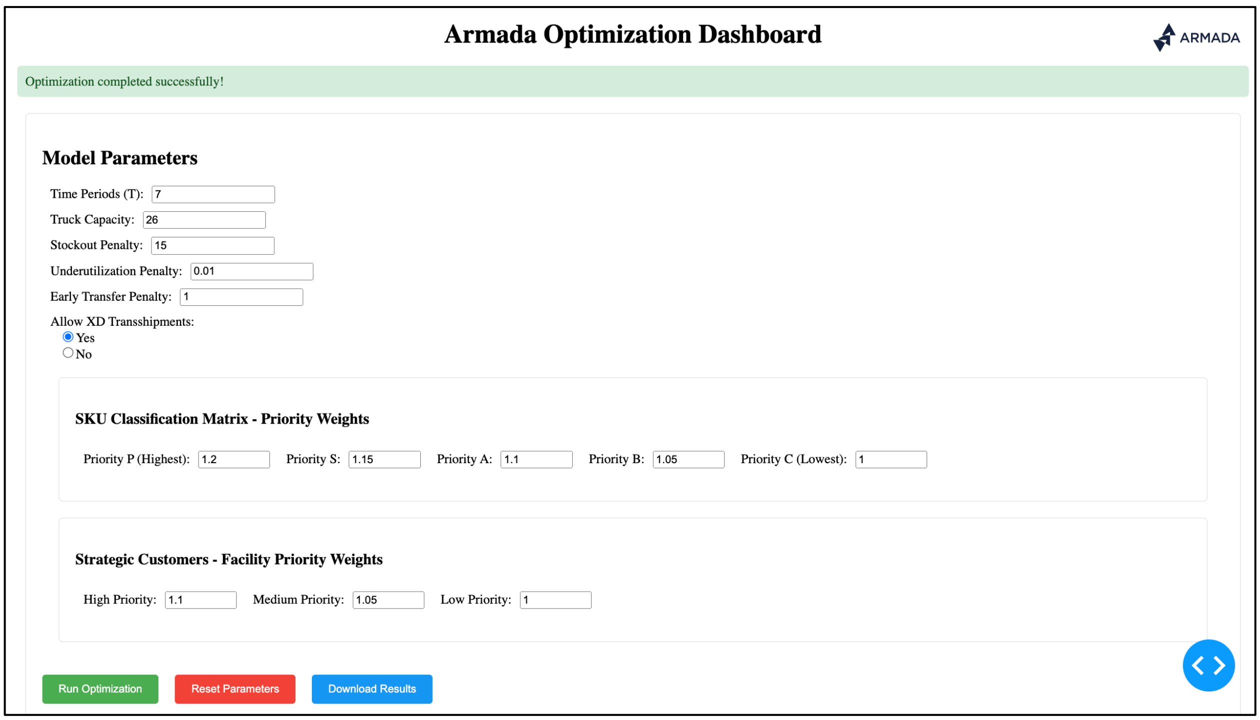

These visual outputs help planners interpret model behavior by showing the distribution of pallet shipments across Dry, Cooler, and Frozen temperature categories, alongside a time-phased view of cost drivers such as stockout penalties, early transfers, and underutilized trucks. This enables quick identification of high cost periods and supports more informed adjustments to replenishment strategy and parameter tuning.

Figure C.2: Dashboard visualizations summarizing pallet movements by temperature type and cost by period

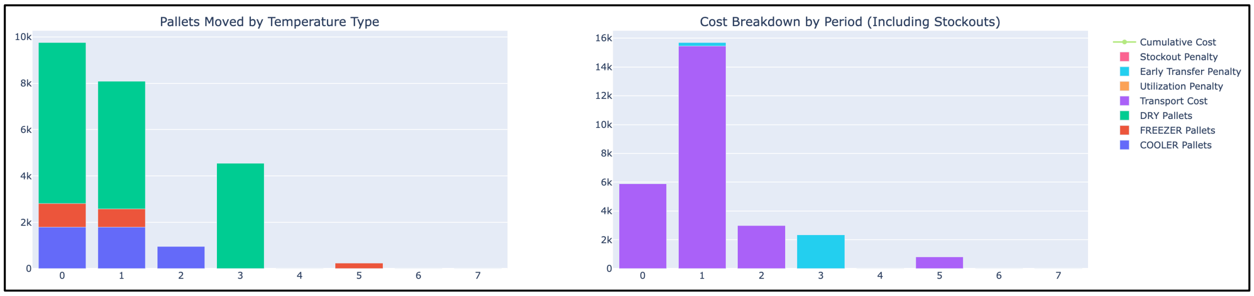

This chart provides a high-level summary of demand recovery outcomes, showing the proportion of total demand that has been successfully covered through planned shipments. In this scenario, the model achieves 100% recovery with no remaining demand at risk, offering planners a clear, visual confirmation that stockout risk has been fully mitigated.

Figure C.3: Dashboard visualization of stockout recovery performance across the network

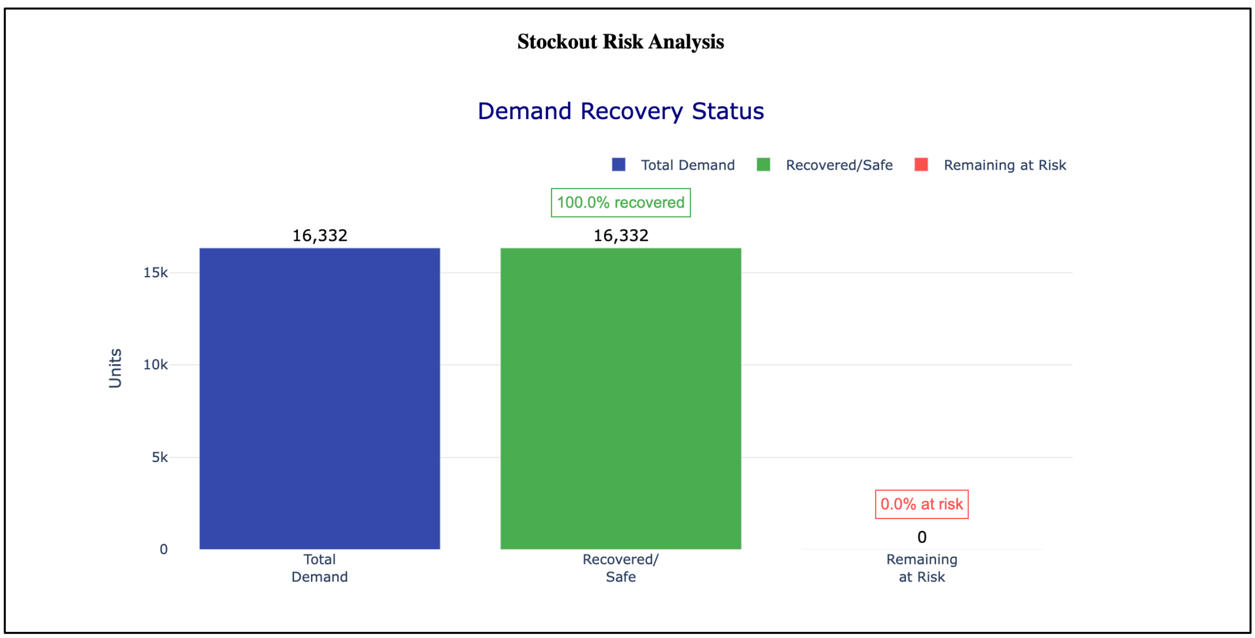

A Shipment Planning Matrix containing sponsor-specific data and a Folium map visualizing the warehousing network, both present in the dashboard, have been excluded from this report to maintain confidentiality and protect sensitive location information.