Aligning Inventory Levels and Distribution Frequencies: Balance Stock Availability With Cost-Effective Distribution Strategies

by

Sayda Elka Benítez Villarreal
Bachelor of Science in Industrial and Systems Engineering, Instituto Tecnológico y de Estudios Superiores de Monterrey, 2019

and

Bryan Lendzion
Bachelor of Arts in Supply Chain Management, Michigan State University, 2022

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

© 2025 Sayda Elka Benítez Villarreal and Bryan Lendzion. All rights reserved.

The authors hereby grant to MIT permission to reproduce and to distribute publicly paper and electronic copies of this capstone document in whole or in part in any medium now known or hereafter created.

Signature of Author: _____
Sayda Elka Benítez Villarreal
Department of Supply Chain Management
May 9, 2025

Signature of Author: _____
Bryan Lendzion
Department of Supply Chain Management
May 9, 2025

Certified by: _____
Dr. Mauricio Gamez Alban
Postdoctoral Associate, MIT Center for Transportation and Logistics
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

Aligning Inventory Levels and Distribution Frequencies: Balance stock availability with cost-effective distribution strategies

by

Sayda Elka Benítez Villarreal

and

Bryan Lendzion

Submitted to the Program in Supply Chain Management

on May 9, 20205 in Partial Fulfillment of the

Requirements for the Degree of Master of Applied Science in Supply Chain Management

ABSTRACT

This capstone project presents a comprehensive inventory management and replenishment strategy developed for Cordis, a U.S.-based manufacturer and distributor of cardiovascular equipment. The project focused on optimizing inventory policies at the Global Distribution and Replenishment Center and the Logistics Service Provider (LSP) for the Europe, Middle East, and Africa (EMEA) region. A tailored methodology was developed using demand segmentation, safety stock calculations, and a combination of continuous and periodic review policies. To evaluate performance, the team modeled eight replenishment scenarios under two safety stock strategies: a proposed formula-based approach and Cordis's current five-weeks-on-hand policy.

Among these, Scenario 8 emerged as the optimal strategy, applying a continuous review policy for high-demand SKUs and periodic replenishment cycles of six and four days for medium- and low-demand SKUs, respectively. This configuration achieved a 35% reduction in operational costs compared to the baseline, while maintaining service level targets. Simulation and inventory analyses further validated that the proposed approach effectively balances cost efficiency with supply chain responsiveness. The results offer Cordis a data-driven, scalable framework to reduce inventory costs, improve performance, and support long-term supply chain resilience.

Capstone Advisor: Dr. Mauricio Gamez Alban
Title: Aligning Inventory Levels and Distribution Frequencies: Balance Stock Availability With Cost-Effective Distribution Strategies

ACKNOWLEDGMENTS

We would like to extend our deepest gratitude to our capstone advisor, Dr. Mauricio Gamez Alban, for his invaluable guidance, support, and encouragement throughout the course of this project. His expertise and mentorship sharpened our analytical thinking and continually challenged us to push beyond conventional boundaries, fostering an environment of learning, growth, and innovation.

Additionally, we sincerely thank our sponsor company, Cordis, and its leadership team for their collaboration, transparency, and trust. Their time, input, and openness made this project possible. We hope that our work contributes meaningful insights and supports their continued pursuit of operational excellence and innovation.

We are also thankful to our writing coach, Ms. Toby Gooley, for her steady support and insightful feedback. Her guidance played a key role in helping us structure our ideas clearly and communicate our research effectively and professionally.

We would like to express our heartfelt appreciation to our families and significant others for their unwavering support, patience, and encouragement throughout these challenging months. Despite the distance, long hours, and demanding schedules, their love and understanding kept us grounded, motivated, and resilient. Finally, we recognize the commitment and collaboration between the two of us—master's students who, through countless meetings, shared knowledge, challenged each other, and offered mutual support. This achievement is a reflection of our joint effort, persistence, and shared dedication.

Table of Contents

1. INTRODUCTION 5

1.1 Background 5

1.2 Motivation 6

1.3 Problem Statement 7

1.4 Scope: Project Goals And Expected Outcomes 7

2. STATE OF THE PRACTICE 8

2.1. Current Inventory Challenges In The Medical Device Industry 8

2.2. Safety Stock 9

2.2. Replenishment Strategies 10

3. METHODOLOGY 12

3.1. Data Cleaning and Data Analysis 13

3.2. Safety Stock Strategy 13

3.3. Inventory Review Policy 14

3.4. Optimization Model 15

4. RESULTS & DISCUSSION 19

4.1. Scenario Analysis 19

4.2. Scenario Results 21

4.3. Inventory Analysis 24

4.4. Proposed Replenishment Strategies 26

4.5. Discussion 29

4.6. Limitations 30

4.7. Managerial Insights 31

5. CONCLUSION 31

5.1 Future Research 32

1. INTRODUCTION

1.1 Background

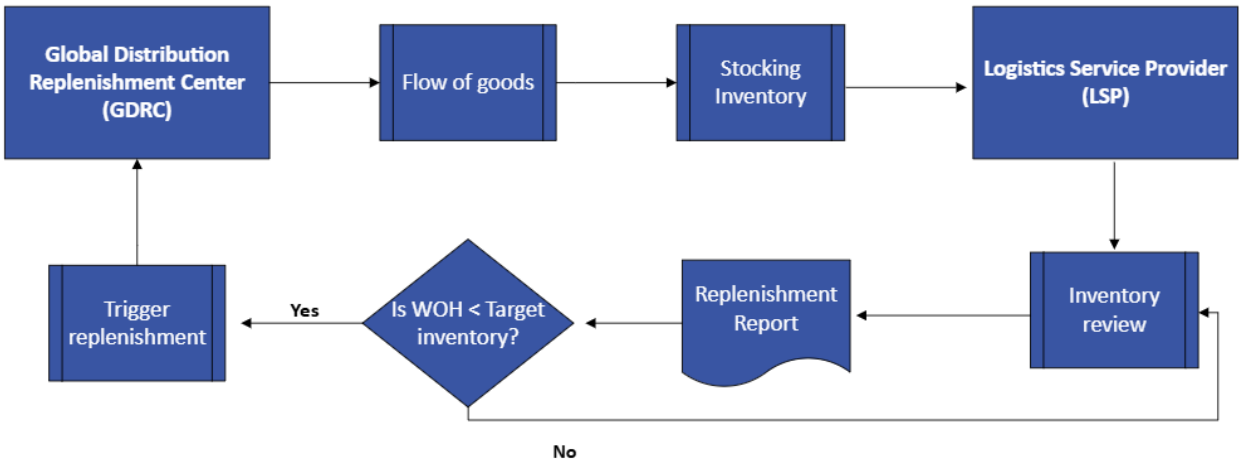
Inventory management and replenishment strategies are critical industry challenges. Companies must evaluate multiple factors when determining the optimal inventory review policy. Demand uncertainty and changes in customer behavior create hurdles, as companies must balance overstocking and understocking inventory. Differing demand patterns across stock-keeping units (SKUs) add complexity, forcing companies to balance ordering frequency and inventory stocking levels. Companies must carefully manage these dynamics to maintain customer service levels while minimizing inventory costs.

Our partner company, Cordis, a U.S.-based manufacturer and distributor of cardiovascular equipment, has faced the challenge of determining effective inventory allocation policies to address demand uncertainty. Their main goal has been to achieve an optimal balance between order frequency and inventory management while meeting customer service targets. The company has adopted an inventory management methodology that emphasizes more frequent ordering to decrease inventory levels for the replenishment of its products. This approach, called "Breakthrough Lean" (BTL), effectively aligns current demand with the required inventory levels, ensuring appropriate replenishment across supply chain nodes.

With the BTL implementation, the company's replenishment processes shifted from a forecast push-based system to a demand pull-based system. Additionally, this approach allowed Cordis to adjust orders more dynamically in response to demand while reducing their backorders from $30 million to $1.2 million in 10 months. The strategy has enabled the company to achieve "just-in-time" (JIT) inventory management, allowing for swift responses to changing demand without the burden of excess stock at each of their seven Logistics Service Providers (LSPs) worldwide.

Figure 1 illustrates the current replenishment process to understand the flow of goods among Cordis' supply chain nodes. Each LSP submits a daily replenishment report and sends a demand signal to order new inventory when its Weeks on Hand (WOH) inventory dips below the target level. The request is sent to the Global Distribution and Replenishment Center (GDRC), located in Mississippi, U.S., where new orders are allocated to its in-stock inventory.

Figure 1. Cordis Supply Chain Replenishment Process from Global Distribution and Replenishment Center to Logistics Service Providers



1.2 Motivation

After implementing the BTL methodology and the demand-pull system, the company established new inventory levels for SKUs aimed at maintaining four to six weeks of stock at their LSPs. This adjustment was defined by the average daily demand for each product over the previous 90 days of its replenishment. Additionally, orders for each LSP now adhere to a new replenishment policy that triggers requirements every two days, allowing the company to respond more effectively to demand uncertainty.

These new strategies have helped improve the company's line-item fill rate (LIFR) percentage while lowering its inventory levels. However, having a dynamic reorder point and a higher inventory order frequency has created extra ordering and material handling costs for the sponsor company. Specifically, the main issues are:

1. Increase in pick and pack time: Workers must handle packages more frequently, especially for low-demand items.
2. Risk of damaging items: Picking smaller quantities requires more package handling and increases the likelihood of product damage.
3. Shipping container capacity issues: The two-day replenishment policy restricts the consolidation of orders for the same SKU, resulting in unused or wasted space within shipping container capacity.
4. Increase in ordering cost: The implementation of dynamic inventory target levels has led to an escalation in ordering costs, as more frequent orders require additional replenishment efforts.

5. Undefined safety stock: The generalized number of Weeks on Hand inventory has resulted in stockouts for high-performing SKUs during unexpected demand spikes, while also leading to higher holding costs for lower-selling SKUs.

These operational challenges have resulted in business issues affecting profitability and customer service. Increased labor demands and rising ordering costs have elevated operational expenses, while stockouts of high-performing SKUs and fulfillment delays threaten customer satisfaction. Additionally, freight density issues and inefficient material handling hinder cost-saving opportunities. Addressing these problems could lower ordering costs, improve service reliability, and define the correct replenishment frequency and inventory levels for each SKU.

1.3 Problem Statement

This capstone project will center on the Global Distribution and Replenishment Center and the top sales LSP for the Europe, Middle East, and Africa region (EMEA LSP). To help the company address the challenges that motivated this project, the primary objective is to establish a methodology for determining the most effective allocation policy and optimizing inventory levels across the supply chain nodes. Accordingly, this capstone aims to provide the sponsor with the optimal deployment frequency and the appropriate safety stock inventory levels for each of the nine selected products at their EMEA LSP.

In that regard, the questions that need to be addressed are:

1. What is the right balance between stock levels and optimal distribution quantities?
2. What is the correct stock level for each selected product?
3. What is the minimum shipment quantity and frequency for each selected product while maintaining the target service level and minimizing costs?

Cordis will provide historical demand, inventory levels, expected forecast, and operational costs for nine selected stock-keeping units (SKUs). Each SKU was selected by the company, using the inventory segmentation technique “ABC inventory categorization” to represent a high-, medium-, and low-demand SKU segment.

1.4 Scope: Project Goals And Expected Outcomes

Our work will encompass an analysis and optimization of the sponsor company's inventory management and distribution processes for their replenishment strategies of nine selected products. We will focus on developing an optimal allocation policy that balances inventory levels at the GDRC and LSP while determining appropriate minimum shipment quantities and frequencies for each ABC inventory categorization product.

To achieve this, we will undertake the following key activities:

1. Data cleaning and analysis: Analyze the company's replenishment report to identify patterns and remove anomalies to optimize inventory planning.
2. Safety stock strategy development: Calculate optimal safety stock inventory for each of the categorized products at the LSP EMEA.
3. Inventory management review policy formulation: Develop continuous and periodic review policies using historical replenishment data to analyze reorder points and total costs.
4. Optimization model creation: Develop an optimization model using programming software to analyze replenishment data, refine it statistically, and determine optimal inventory parameters to minimize costs and meet operational constraints.

The primary deliverables to the sponsor company will include:

1. A report detailing our analysis, findings, and recommendations.
2. An optimized inventory allocation policy tailored to their specific product categories.
3. A set of guidelines for minimum shipment quantities and frequencies for each ABC inventory category.
4. A customized inventory management tool or model that the company can use for ongoing optimization.

These deliverables will enable the sponsor company to implement more efficient distribution processes, reduce small shipment quantities, and achieve an optimal balance in inventory levels across their supply chain. The expected outcomes include reduced operational costs, improved service levels, and enhanced overall supply chain efficiency.

2. STATE OF THE PRACTICE

The primary challenge of our capstone is how our sponsor company can determine the optimal replenishment frequency and ideal safety stock levels to reduce costs related to the ordering process for its LSPs. To address this, we started by examining the existing inventory management challenges within the medical device industry. Next, we explored widely adopted strategies for assessing inventory levels. Finally, we identified various replenishment policy methods to establish the most effective order frequency.

2.1. Current Inventory Challenges In The Medical Device Industry

Inventory management in the healthcare industry is complex and costly, particularly for medical devices. Supply chain-related expenses represent a substantial portion of hospital costs, and this financial

pressure is often transferred to medical device suppliers. Consignment policies further complicate this, where hospitals are charged only once as they use the medical devices. Alternatively, supply chain optimization offers many cost-saving opportunities. Research presented to the Medical Device Supply Chain Council estimated that supply chain expenses account for 40–45% of hospital operating costs, and the same study found that enhanced supply chains could potentially reduce these costs by 5–15% (Global Healthcare Exchange, 2011).

As the global demographic trend shifts towards an aging population, the healthcare and medical devices industries are encountering significant challenges within their supply chains. The increasing demand for medical devices, coupled with stringent production regulations, underscores the vital role of distributors, as the LSPs mentioned in Section 1.1, in maintaining customer satisfaction and mitigating supply shortages. Their effective management is essential to navigating these complexities and ensuring that the needs of this growing demographic are met (Anjani & Nizar, 2021).

Moreover, a distinct challenge arises for items without clearly defined inventory levels correlated with their sales performance. Hafnika et al. (2016) conducted a case study to determine optimal safety stock levels within this field, developing a probabilistic inventory model that considered variables such as lead time and fluctuating demand. Additionally, by implementing a continuous review policy, which triggers orders as inventory approaches its reorder point, they successfully reduced overstock costs in the pharmacy unit by 57%.

In the medical device industry, maintaining high inventory levels is a common strategy to address urgent customer demands. However, this approach frequently results in excessive stock, especially within supply chains characterized by long lead times and items such as sterilized surgical products with limited shelf life. By using the Strategic Inventory Placement optimization method, Tjhin & Pandey (2012) led to a reduction of over 60% in inventory value and a decrease of more than 20% in operating costs. Additionally, overall supply chain efficiency was significantly enhanced by prioritizing finished product inventory and minimizing component stock.

2.2. Safety Stock

Safety stock calculation is a widely used inventory management strategy to mitigate risks associated with demand variability and lead time uncertainty. Maintaining an additional buffer of inventory ensures that fluctuations in actual demand, which may differ from forecasts, do not lead to stockouts. Safety stock also accounts for discrepancies between expected and actual replenishment times, helping to maintain smooth operations and prevent supply disruptions.

There are three main categories for calculating appropriate safety stock levels for individual items: safety stocks derived from a basic or simplistic approach, safety stocks determined by customer service levels, and safety stocks aimed at minimizing estimated shortage and inventory holding costs (Johnsson & Mattsson, 2019).

Hillstrom and Malabanan (2013) identified the critical trade-off between holding costs and stockout inventory costs. Factors such as accuracy in customer service levels, lead time replenishment, expected demand, standard costs, and stockout probabilities offer valuable insights for reducing inventory levels. Furthermore, they emphasized the importance of product categorization in addressing items appropriately. For fast-moving SKUs, relying solely on cycle stock can effectively meet service targets, while slow-moving items require a limited amount of safety stock to ensure adequate availability.

By maintaining an appropriate buffer, businesses can safeguard against stockouts and ensure operational continuity. The insights from both Johnsson and Mattsson (2019) and Hillstrom and Malabanan (2013) underline the importance of a strategic approach in determining safety stock levels, which must consider customer service requirements, cost implications, and product categorization. Ultimately, a well-calibrated safety stock strategy enhances service levels and optimizes inventory holdings.

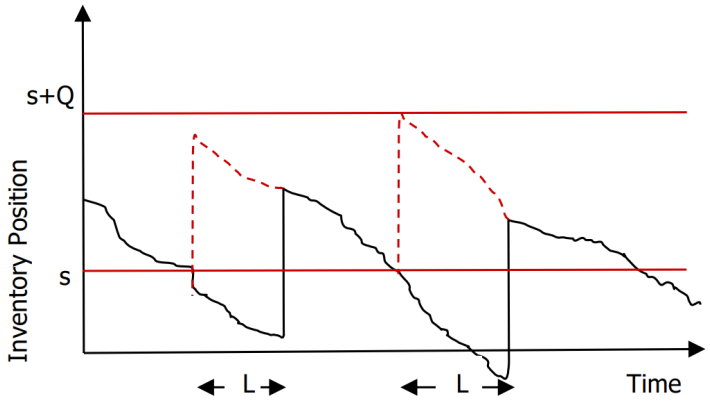
2.2. Replenishment Strategies

To minimize costs while maintaining a high service level, it is necessary to find the optimal inventory management policy for each ABC-classified SKU. Two stochastic demand replenishment policies are continuous review and periodic review. These policies are not comprehensive; however, they can be strategically used for different SKU segmentation to optimize service levels while minimizing total relevant costs (Setyaningsih & Basri, 2013).

As illustrated in Figure 2, the continuous review method employs a reorder point (ROP) to monitor inventory levels. When inventory reaches or falls below this predetermined point, orders are initiated, usually involving fixed quantities such as the Economic Order Quantity (EOQ). This strategy is especially advantageous for high-value or critical items, where stockouts can have severe consequences.

The primary benefits include a reduction in stockouts, the maintenance of optimal inventory levels, and its effectiveness in managing volatile or unpredictable demand scenarios. This enhances product visibility and ensures accurate accounting calculations by providing up-to-date data on the cost of goods sold. However, its primary limitation lies in the high implementation cost, as it requires constant maintenance and is more susceptible to errors.

Figure 2. Continuous Review: Order-Point, Order-Quantity (s,Q) (Correll, 2022)



In contrast, the periodic review policy involves assessing inventory at predetermined intervals, such as daily, weekly, or monthly. If inventory levels fall below the target threshold during these reviews, a replenishment order is issued, as described in Figure 3, with the parameters described in Table 1. This method is cost-effective, as it minimizes the need for real-time monitoring and complex tracking systems. It is especially suitable for items with stable and predictable demand, as well as low-value products like office supplies and basic commodities.

Key advantages of this policy include its ease of implementation, cost savings, and the ability to optimize scheduling for both suppliers and inventory management. In contrast, its disadvantages include potential inaccuracies in high-volume environments and the need for managers to make assumptions about inventory and supply levels during the review period (Setyaningsih & Basri, 2013).

Figure 3. Periodic Review: Order-Up-To-Level (R, S) (Correll, 2022)

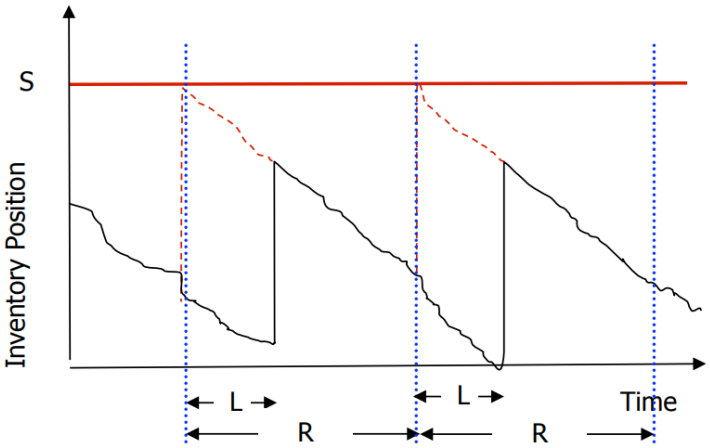


Table 1. Review Policy Notations

Notation	*Definition*
s	Reorder Point
Q	Order Quantity
S	Order-up-to Level
R	Review Period
L	Replenishment Lead Time
IOH	Inventory On Hand
IP	Inventory Position = (IOH) + (Inventory On Order) − (Backorders)

The literature review on current inventory challenges in the medical device industry, safety stock, and replenishment policies laid a strong foundation for developing the methodology for this capstone project. Through an in-depth examination of existing research, several industry-specific challenges emerged, including demand variability, stringent regulatory requirements, and disruptions within the supply chain. These insights underscore the pressing need for robust inventory management strategies tailored to the unique circumstances of the medical device sector.

3. METHODOLOGY

In this study, we developed a methodology to establish effective inventory management policies for the sponsor company. Our approach, grounded in an extensive literature review, integrates several key strategies tailored to the specific needs and market conditions of the company. The methodology encompasses the development of customized inventory review policies for each SKU segment, recognizing that different product categories may require distinct management approaches. We set safety stock levels based on individual SKU demand performance, assuming a constant lead time and targeting an ambitious 98% cycle service level to ensure high customer satisfaction.

To validate and refine our approach, we conducted rigorous scenario analyses, enabling us to determine the most appropriate inventory review policy for each category under various potential market conditions. Based on this analysis, we developed inventory management strategies, including minimum order quantities for efficient replenishment. Crucially, our methodology acknowledges the variability in SKU demand patterns, categorizing items into high-medium-low demand segments. This segmentation allows for the application of specific inventory review policies and safety stock levels tailored to each demand category, ensuring a balanced approach that optimizes inventory levels while maintaining service quality across the entire product range.

3.1. Data Cleaning and Data Analysis

Data cleaning is vital for ensuring accuracy and reliability by removing errors and irrelevant information that could lead to misleading conclusions. Data analysis then converts this cleaned data into actionable insights, enabling informed decision-making and trend identification. Therefore, these two processes are implemented as the initial steps in the methodology applied to the information provided by Cordis.

The data provided by the company includes two primary reports covering six months: the daily replenishment report and the orders processed report. The daily replenishment report generated by the LSP (in this case, EMEA) offers SKU-level insights into on-hand and in-transit inventory, average daily demand, target inventory levels, and the quantity needed to meet those targets. Meanwhile, the orders processed report details the number of orders processed and shipped by the GDRC at the SKU level, including the quantity shipped per SKU for each order and the corresponding request date.

By employing descriptive statistics, including measures of central tendency such as mean and median, as well as variability metrics like standard deviation and range, and frequency histograms and boxplots, we identified the key parameters and characteristic patterns of demand and inventory variation for each of the nine SKUs. This analysis facilitated the detection and removal of anomalies within the data. Furthermore, it identified trends in demand fluctuations that could impact inventory planning. These insights provide a solid foundation for optimizing stock management and enhancing data accuracy for future analyses.

3.2. Safety Stock Strategy

Considering that the company has a processing plus transit lead time to the Logistics Service Provider (LSP) in EMEA, Netherlands, of 14 days on average, as well as a target cycle service level (CSL) of 98%, the safety stock was determined using the formula provided by Silver et al. (1998), shown in Equation 1.

$$Safety\ Stock\ = k * \sigma * \sqrt{L + R} \quad \text{(Eq.1)}$$

Where k represents the safety stock factor corresponding to the CSL, σ is the standard deviation of the demand, L is the total lead time, and R denotes the review period determined by the inventory review policy. This approach reduces the risk of stockouts, improves customer satisfaction by ensuring product availability, and helps prevent overstocking, which leads to unnecessary holding costs.

Additionally, to assess the minimum safety stock required at the EMEA LSP, a lower bound was established based on item performance. This lower bound assumes a minimum safety stock equivalent to

demand over half the average lead time, providing a baseline for comparison and ensuring a minimum level of buffer inventory.

3.3. Inventory Review Policy

SKU demand (high, medium, or low) determines the appropriate inventory review policy and safety stock level. When considering the continuous review policy (s,Q), we used equation 2 formula to find the Economic Order Quantity (EOQ) and equation 3 formula to find the reorder point (s). Both formulas were provided by Correll (2022), and tables 2 and 3 explain each parameter, respectively.

$$Q^* = \sqrt{\frac{2c_t D}{c_e}} \quad \text{(Eq.2)}$$

Table 2. Parameters for Economic Order Quantity, Continuous Review policy

Parameter	*Definition*
Q^*	Optimal order quantity minimizing total inventory costs (ordering + holding)
c_t	Fixed cost per order (e.g., administrative or shipping costs)
D	Annual demand (total yearly quantity required)
c_e	Annual holding costs per unit (e.g., storage, insurance, opportunity cost)

Table 3. Parameters for Reorder Point, Continuous Review policy

Parameter	*Definition*
s	Inventory level to trigger a replenishment order
μDL	Average demand during lead time
σDL	Standard deviation of demand during lead time
k	Safety stock factor based on the desired service level

$$s = \mu DL + k * \sigma DL \quad \text{(Eq.3)}$$

In the (s, Q) continuous review policy, the Economic Order Quantity (EOQ) model is used to determine the optimal order quantity (Q^*) that minimizes the total cost of ordering and holding inventory. The reorder point (s) indicates the inventory level at which a new order should be placed. In practice, this means ordering Q units whenever the inventory position falls to or below s.

In contrast, the (R, S) periodic review policy sets a fixed review interval (R) and calculates a target inventory level (S) using Equation 4 with its parameter description in Table 4 to ensure adequate stock throughout each review cycle. As highlighted by Correll (2022), selecting the appropriate inventory review policy requires alignment with the characteristics of each SKU to ensure cost-efficient inventory control.

$$S = \mu_{DL+R} + k * \sigma_{DL+R} \quad \text{(Eq.4)}$$

Table 4. Parameters for Reorder Point, Periodic Review policy

Parameter	*Definition*
S	Target inventory level
$\mu DL + R$	Average demand over the combined review period and lead time
σDL	Standard deviation of demand during lead time
k	Safety stock factor based on the desired service level

3.4. Optimization Model

After conducting a thorough inventory analysis using standard inventory review formulas that best fit the company's current circumstances, the next step was to develop an optimization model to better capture the frequency, replenishment cycles, inventory and shipment quantities for each of the company's products. Our optimization model was coded in C++ and solved using specialized programming software such as CPLEX. This model will serve as a tool for identifying the most effective inventory strategy for each product in the portfolio.

The optimization model is designed to extract critical parameters by utilizing data derived from replenishment reports for final decision-making. The process begins with a descriptive statistical analysis of the available data, which provides insights into trends and variabilities. Following this, we cleaned the data by removing outliers that could skew results and established coherent patterns for final demand and order quantities.

After data cleaning, the model employs quantitative techniques to determine optimal safety stock levels, order points, and replenishment frequencies. These parameters are strategically crafted to minimize total costs while accommodating operational constraints, such as a 14-day lead time from the company's Global Distribution Replenishment Center to the LSP in Europe.

This model aims to determine the optimal replenishment quantity and frequency for each product classification, minimizing associated costs. The model consists of three key sets defined based on the data provided. Let i represent the set of stock-keeping units (SKUs), where i = {0,1,…,8}, corresponding to the nine selected items. Set j denotes the facility locations, where j = {0,1,2}, which includes the production factory, the GDRC, and the logistics service provider. Finally, t represents the set of time periods, spanning a 139-day timeline during which operations are analyzed.

Table 5 lists the parameters used in the formulation presented below based on the defined sets. These parameters play a critical role in accurately capturing the dynamics of the supply chain and ensuring the robustness of the optimization model. as well as the safety stock calculation, using the formula from the method mentioned in Section 3.2.

Table 5. Parameters selected for the optimization model

Parameter	*Definition*
cah_j	Receiving capacity at facility j
cas_j	Shipping capacity at facility j
d_{it}	Demand for product i at time t
lt_j	Travel (lead) time from facility j to the next facility
pco_i	Safety stock penalty cost for product i
sco_j	Storage (holding inventory) cost at facility j
st_{ij}	Safety stock goal for product i at facility j
tco_{ij}	Processing cost for product i from facility j to the next facility
iv_{ij}	Initial inventory for product i at facility j
stv_j	Backorder cost at facility j
$Cycle_i$	Replenishment cycle for product i
lb_{ij}	Safety stock lower bound for product i at facility j
M	A large constant used to activate or deactivate constraints
v	Capacity per pallet shipped
m	Cost per pallet shipped

As decision variables of this problem, x_{ijt} represents the quantity of goods shipped of the product type i at facility j at time t; y_{ijt} is the amount of inventory of the product type i from facility j at time t; $b_{i\,t}$ number of back orders for of SKU i at time t; s_{ijt} amount of inventory below their safety stock of the product type i at facility j at time t; o_{jt} Pallets shipped from the facility j at time t; and w_{ijt} binary variable to replenish the product i at facility j during time t.

The mathematical model for the proposed problem is formulated as follows for its objective function to minimize cost, as shown in Equation 5. Moreover, equations 6 through 10 represent the components for the main function, where z_1 represents the total processing cost for each quantity shipped, z_2 denotes the total inventory holding, z_3 and z_4 show the total penalty cost for back orders and non-accomplished safety stock level, respectively, and z_5 the cost per pallet shipped.

$$Min\ z = z_1 + z_2 + z_3 + z_4 + z_5 \quad \text{(Eq.5)}$$

$$z_1 = \sum_{i\in I}\sum_{j\in J}\sum_{t\in T} tco_{ij} * x_{ijt} \quad \text{(Eq.6)}$$

$$z_2 = \sum_{i\in I}\sum_{j\in J}\sum_{t\in T} sco_j * y_{ijt} \quad \text{(Eq.7)}$$

$$z_3 = \sum_{j\in J}\sum_{t\in T} stv_j * b_{it} \quad \text{(Eq.8)}$$

$$z_4 = \sum_{i\in I}\sum_{j\in J}\sum_{t\in T} pco_i * s_{ijt} \quad \text{(Eq.9)}$$

$$z_5 = \sum_{j \in J} \sum_{t \in T} o_{jt} * m \quad \text{(Eq.10)}$$

Additionally, the model incorporates operational constraints to ensure efficient merchandise flow. Equations 11 through 13 address demand satisfaction by accounting for current and back orders from the previous day, ensuring that the quantity shipped from the preceding facility fulfills requirements. Regarding inventory in the facilities (1 and 2), equations 14 through 17 refer to inventory management in the LSP EMEA facility. These equations consider inventory from the previous day, incoming inventory from the previous facility, and outgoing inventory designated for customers. Similarly, equations 18 through 20 address inventory management for the GDRC, following the same principles. Additionally, equations 21 and 22 represent the constraints to make the safety stock at EMEA satisfy both the goal and the lower bound.

Equations 23 and 24 detail constraints addressing each facility's receiving and shipping capacities based on the quantity shipped. Meanwhile, equations 25 and 26 establish the relationship constraint to ensure that the inventory before the lead time at each facility is received on the first available day following the 14-day lead time for the LSP and GDRC. Equation 27 specifies that the decision variable for pallets shipped aligns with the company's assumption that each pallet has a maximum capacity of 400 boxes.

Moreover, equations 28 and 29 define the logic for controlling replenishment timing. The first equation uses the Big M method to allow replenishment only when a binary decision variable is active; otherwise, no quantity can be assigned. The second equation ensures that replenishments are not scheduled for consecutive periods, enforcing the minimum time between orders based on the defined replenishment cycle for each SKU and location. Lastly, equations 30 through 34 define the non-negativity constraint, always ensuring positive values for the decision variables, as well as the binary variable in equation 35.

$$x_{i2(t-lt_2)} + b_{it} = d_{it} \qquad \forall\, i \in I, t \in T \mid t = lt_2 \qquad \text{(Eq.11)}$$

$$b_{it} = 0 \qquad \forall\, i \in I, t \in T \mid t < lt_2 \qquad \text{(Eq.12)}$$

$$x_{i2(t-lt_2)} + b_{it} = d_{it} + b_{i(t-1)} \qquad \forall\, i \in I, t \in T \mid t > lt_2 \qquad \text{(Eq.13)}$$

$$y_{i2t} = y_{i2(t-1)} - x_{i2t} \qquad \forall\, i \in I, t \in T \mid 1 \leq t < lt_1 \qquad \text{(Eq.14)}$$

$$y_{i2t} = y_{i2(t-1)} + x_{i1(t-lt_1)} - x_{i2t} \qquad \forall\, i \in I, t \in T \mid lt_1 > 1, t \geq lt_1 \qquad \text{(Eq.15)}$$

$$y_{i20} = iv_{i2} - x_{i20} \qquad \forall\, i \in I \qquad \text{(Eq.16)}$$

$$y_{i2t} = y_{i2(t-1)} + x_{i1(t-lt_1)} - x_{i2t} \qquad \forall\, i \in I, t \in T \mid lt_1 \leq 11, t \geq 1 \qquad \text{(Eq.17)}$$

$$y_{i1t} = iv_{i1} - x_{i10} \qquad \forall\, i \in I, t \in T \mid t < lt_0 \qquad \text{(Eq.18)}$$

$$y_{i1t} = y_{i1(t-1)} + x_{i0(t-lt_0)} - x_{i1t} \qquad \forall\, i \in I, t \in T \mid t \geq lt_0, lt_0 > 1 \qquad \text{(Eq.19)}$$

$$y_{i1t} = y_{i1(t-1)} + x_{i0(t-lt_0)} - x_{i1t} \qquad \forall\, i \in I, t \in T \mid t \geq 1, lt_0 \leq 1 \qquad \text{(Eq.20)}$$

$$y_{i2t} + s_{i2t} \geq st_{i2} \qquad \forall\, i \in I, t \in T \mid t \geq lt_0 + lt_1 \qquad \text{(Eq.21)}$$

$$y_{i2t} \geq lb_{i2} \qquad \forall\, i \in I, t \in T \mid t \geq lt_0 + lt_1 \qquad \text{(Eq.22)}$$

$$\sum_{i \in I} x_{ijt} \leq cah_{(j+1)} \qquad \forall\, j = \{0, .., |J| - 1\}, t \in T \qquad \text{(Eq.23)}$$

$$\sum_{i \in I} x_{ijt} \leq cas_j \qquad \forall\, j \in J, t \in T \qquad \text{(Eq.24)}$$

$$x_{i1t} = 0 \qquad \forall\, i \in I, t \in T \mid t < lt_0 \qquad \text{(Eq.25)}$$

$$x_{i2(|T|-1)} = 0 \qquad \forall\, i \in I \qquad \text{(Eq.26)}$$

$$o_{jt} \geq x_{ijt}/v \qquad \forall\, i \in I, j \in J, t \in T \qquad \text{(Eq.27)}$$

$$x_{ijt} \leq M * w_{ijt} \qquad \forall\, i \in I, j \in J, t \in T \qquad \text{(Eq.28)}$$

$$w_{i1t} + w_{i1t'} \leq 1 \qquad \forall\, i \in I, t = \{0, .., |T| - 1\}, t' = \{t + 1, .., |T|\} \mid$$

$$t' = \{t + 1, .., |T|\} \mid t' - t < cycle_i \qquad \text{(Eq.29)}$$

$$x_{ijt} \geq 0 \qquad \forall\, i \in I, j \in J, t \in T \qquad \text{(Eq.30)}$$

$$y_{ijt} \geq 0 \qquad \forall\, i \in I, j \in J, t \in T \qquad \text{(Eq.31)}$$

$$s_{ijt} \geq 0 \qquad \forall\, i \in I, j \in J, t \in T \qquad \text{(Eq.32)}$$

$$b_{it} \geq 0 \qquad \forall\, i \in I, t \in T \qquad \text{(Eq.33)}$$

$$o_{jt} \geq 0 \qquad \forall\, j \in J, t \in T \qquad \text{(Eq.34)}$$

$w_{ijt} \in \{0,1\}$ $\forall \ i \in I, j \in J, t \in T$ (Eq.35)

4. RESULTS & DISCUSSION

This chapter presents the results of the scenario-based analysis conducted to evaluate alternative inventory management strategies for Cordis. The study compared two safety stock approaches—the proposed formula developed in this project and Cordis's current method of maintaining five weeks on hand—across eight replenishment scenarios. Scenario 1 serves as the baseline scenario for both safety stock approaches, reflecting the company's existing practice of applying a continuous review policy with daily replenishment across all SKUs. This setup enables maximum flexibility and responsiveness but often results in higher operational costs due to frequent shipments and handling. Scenario 1 provides a reference point for assessing the performance of the seven alternative strategies.

The remaining seven scenarios were designed to test the impact of differentiated replenishment cycles—particularly for medium- and low-demand SKUs—by applying periodic review policies with shipment intervals of 2, 4, or 6 days. These scenarios aim to consolidate shipments and reduce replenishment frequency, thereby lowering operational burden without compromising service levels. High-demand SKUs in all scenarios continued to follow a continuous review policy to maintain responsiveness.

As detailed in Section 3.4: Optimization Model, operational cost in our analysis includes three key components: (1) inventory holding cost at both the Global Distribution and Replenishment Center (GDRC) and the Logistics Service Provider (LSP), (2) transportation fixed cost, based on a charge per pallet shipped (parameter z5), and (3) transportation variable cost, calculated from the total processing cost for each quantity shipped.

The analysis showed that Scenario 8, in a quarter-of-a-year simulation, achieved the best operational performance under the proposed safety stock model. This scenario applies to a continuous review policy for high-demand SKUs (0-day cycle), and periodic cycles of 6 and 4 days for medium- and low-demand SKUs, respectively. With an average operational cost of USD $1,960,932.18, Scenario 8 demonstrates a 35% cost reduction compared to the baseline. These results underscore the value of aligning replenishment cycles with SKU demand profiles to achieve significant cost savings while maintaining service reliability, offering Cordis a scalable, data-driven alternative to its current one-size-fits-all strategy.

4.1. Scenario Analysis

After carefully selecting the formulas for the review policies and safety stock calculations and thoroughly assessing the cleaned data using descriptive statistics, we developed and analyzed multiple

scenarios for the nine SKUs. The primary objective was to identify the most cost-effective inventory policy that balances service level requirements with operational efficiency.

Initial discussions with the sponsor company and the results from the first model run revealed that due to high demand, Category A SKUs are well-suited to a continuous review policy. This approach aligns with the company's current replenishment strategy, enabling frequent ordering while maintaining high service levels.

To evaluate potential cost savings and assess the impact of a periodic review policy, we created eight alternative scenarios focused on Categories B and C—medium- and low-demand SKUs. In collaboration with the sponsor, we tested replenishment cycles of 0, 2, 4, and 6 days to determine the feasibility of consolidating orders and reducing operational costs.

Table 6 presents the replenishment cycle lengths (in days) applied to each SKU category across eight defined scenarios. A value of 0 indicates a continuous review policy, while numeric values represent periodic review cycles. The key characteristics of each scenario are as follows:

- Category A remains under a continuous review policy (0 days) across all scenarios, reflecting its high and stable demand.
- Scenario 1 represents a continuous review policy for all SKUs. This scenario serves as a baseline, allowing the model to optimize inventory decisions without any constraints on replenishment cycles.
- Scenarios 2–4 consistently use 2-, 4-, and 6-day replenishment cycles, respectively. However, they shift the application of these cycles between Category B (medium demand) and Category C (low demand). This design directly compares how each category responds to the same replenishment cycle when managed under a periodic review policy.
- Scenarios 5 and 6 apply a fixed 2-day replenishment cycle to Category B, while Category C cycles vary (4 and 6 days, respectively).
- Scenarios 7 and 8 assign periodic policies to both Categories B and C simultaneously, using 4-day and 6-day replenishment cycles.

This structure enables the evaluation of cost, inventory levels, and service performance under different policy configurations, helping to determine the most effective inventory review strategy for each category.

Table 6. Replenishment Cycles by Category in Each Defined Scenario

Category	*Scenario 1*	*Scenario 2*	*Scenario 3*	*Scenario 4*	*Scenario 5*	*Scenario 6*	*Scenario 7*	*Scenario 8*
A	0	0	0	0	0	0	0	0
B	0	2	4	6	2	2	4	6
C	0	2	4	6	4	6	6	4

Additionally, we ran the model for each scenario using the safety stock formulas established in Section 3.2. to assess benchmark performance and ensure a meaningful comparison, we also evaluated the model outcomes against the company's current safety stock targets, set at five weeks of average daily demand (ADD) as the standard and two days of ADD as the lower bound. This allowed us to directly compare the proposed methodology to existing practices and assess cost efficiency and service level improvements. The goal was to identify the most cost-efficient replenishment policy and evaluate how different safety stock assumptions impact overall system performance and operational trade-offs.

This comprehensive approach enabled us to identify the most suitable review policy for each product category, determine optimal safety stock levels for individual SKUs, and evaluate the resilience of each policy under varying demand patterns and potential supply chain disruptions. By simulating a range of realistic scenarios, we ensured that the recommended strategies support high customer service levels while minimizing total costs. As a result, we can confidently propose inventory policies that are not only cost-effective but also robust enough to adapt to market volatility and operational uncertainties.

4.2. Scenario Results

Among all eight scenarios tested, Scenario 1 consistently produced the lowest total cost under both safety stock methods, as shown in Table 7. This scenario applies a continuous review policy with daily replenishment for all SKUs, giving the model maximum flexibility in determining optimal order timing. Using the proposed safety stock formula, the total cost in Scenario 1 was USD 2,125,506.79, while under Cordis's current safety stock approach, the total cost reached USD 3,134,081.77. To calculate this value, we ran the optimization model using the company's existing safety stock policy, targeting five weeks of average daily demand as the goal and two weeks as the lower bound, while keeping all other model parameters and data consistent.

As shown in Table 7, we also calculated the percentage reduction in total cost achieved by the proposed strategy across all eight scenarios. The results demonstrate consistent improvements, with cost reductions ranging from 30.8% to 32.3%. This reinforces the value of a data-driven safety stock approach, particularly when paired with optimized replenishment cycles tailored to SKU demand levels. These results highlight the potential cost savings of a dynamic, formula-driven continuous review strategy,

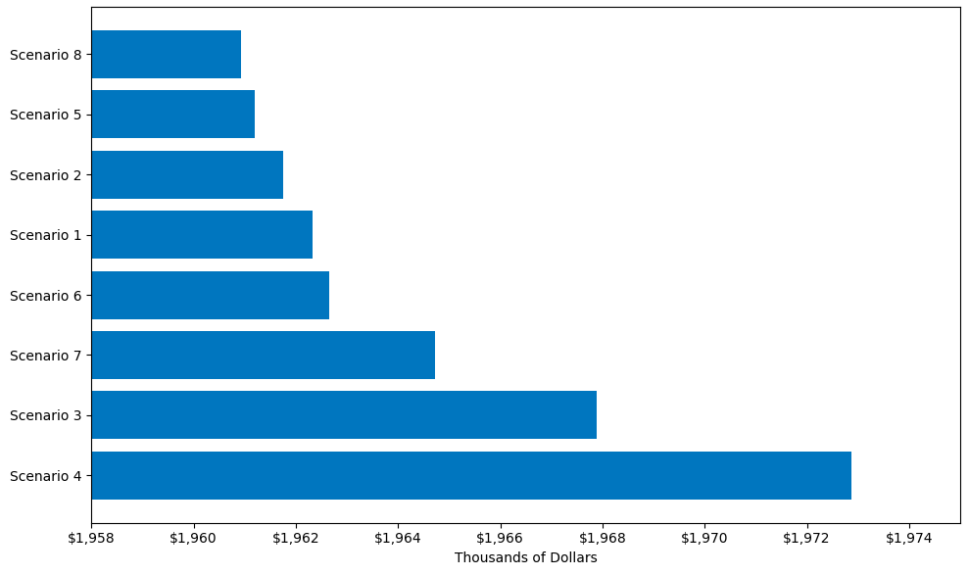
especially for high-demand SKUs. However, it is important to note that while Scenario 1 offers strong cost advantages, its real-world feasibility may be limited by the operational burden of daily replenishment for medium- and low-demand items.

Table7. Average Total Cost Comparison Across Scenarios: Proposed vs. Current Safety Stock Strategy (USD)

Scenario	*Proposed Strategy*	*Current Strategy*	*% Reduction*
Scenario 1	$2,125,506.79	$3,134,081.77	32.3%
Scenario 2	$ 2,127,982.91	$3,140,283.44	32.3%
Scenario 3	$ 2,144,348.16	$3,139,736.82	31.7%
Scenario 4	$2,152,933.13	$3,144,723.77	31.5%
Scenario 5	$2,129,805.99	$3,144,705.17	32.3%
Scenario 6	$2,132,743.53	$3,139,532.14	32.0%
Scenario 7	$2,146,583.79	$3,136,662.08	31.6%
Scenario 8	$2,177,828.72	$3,144,975.73	30.8%

Figure 4 illustrates the operational cost results for each of the eight replenishment scenarios evaluated under the proposed inventory strategy. Each bar represents the total operational cost—covering inventory holding and transportation expenses—associated with a specific replenishment cycle configuration for SKU categories A, B, and C. The graph clearly shows variation in cost performance across scenarios, highlighting the impact of different replenishment cycle combinations on overall efficiency.

Figure 4. Average Operational Cost Across Scenarios Using the Proposed Safety Stock Model

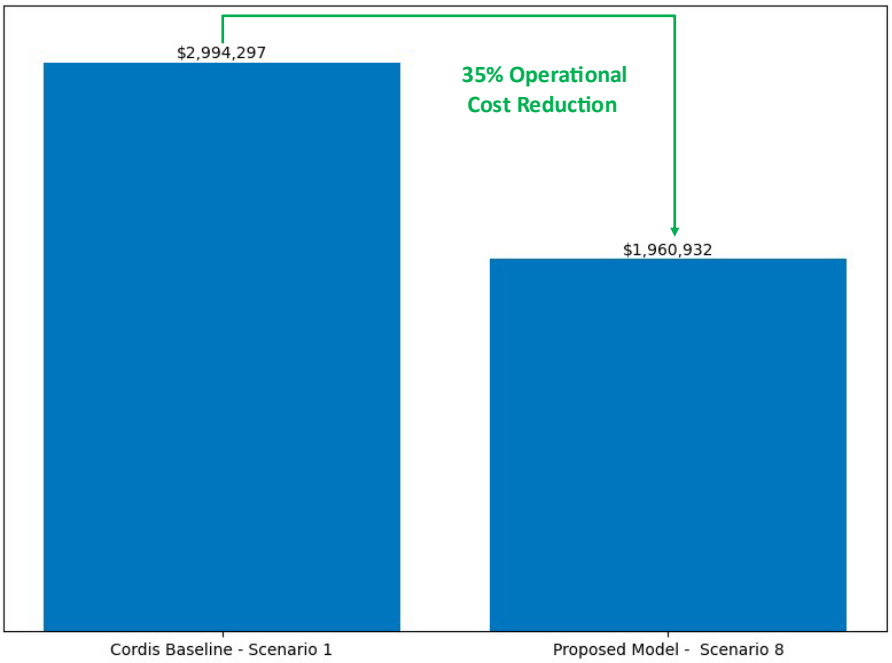


Among the scenarios, Scenario 8 achieved the lowest operational cost, totaling USD 1,960,932, followed closely by Scenarios 5 and 2, with costs of USD 1,961,197 and USD 1,961,757, respectively. At the other end of the spectrum, Scenario 4 exhibited the highest operational cost at USD 1,972,870, indicating that overly infrequent replenishment for Category B and C SKUs may lead to inefficiencies. Other higher-cost scenarios include Scenario 3 and Scenario 7, with operational costs of USD 1,967,883 and USD 1,964,725, respectively.

The figure 4 demonstrates that strategic variation in replenishment cycles—particularly for medium- and low-demand SKUs—can substantially impact operational costs. The more efficient scenarios tend to apply moderate replenishment cycles that balance order frequency and inventory holding, confirming the importance of demand-aligned review policies.

Figure 5 compares operational costs directly between Cordis's current strategy (Scenario 1) and the most cost-effective configuration identified under the proposed model (Scenario 8). Scenario 1, which reflects the company's baseline approach using continuous review and fixed safety stock levels, resulted in an operational cost of USD 2,994,297. In contrast, Scenario 8 achieved a significantly lower price of USD 1,960,932, representing a 35% reduction. This cost difference visually underscores the potential savings of shifting from a uniform, high-frequency replenishment strategy to a more segmented, demand-driven approach.

Figure 5. Operational Cost Analysis: Proposed Solution vs. Cordis Baseline in USD



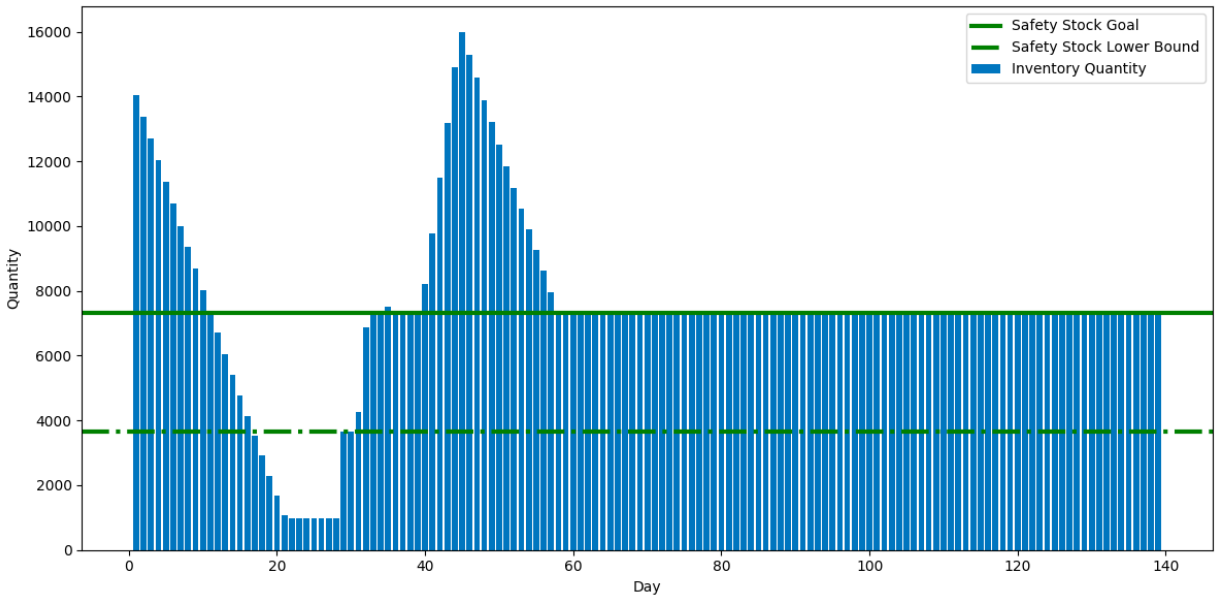
Scenario 8 applies a continuous review policy for high-demand Category A SKUs (0-day cycle) while assigning periodic replenishment cycles of six days for Category B and four days for Category C. This hybrid strategy allows Cordis to reduce unnecessary replenishment frequency for lower-demand items while maintaining responsiveness for critical products. These results highlight the value of aligning replenishment policies with demand characteristics, demonstrating that a tailored strategy can significantly improve cost efficiency without compromising service levels.

4.3. Inventory Analysis

Continuing the analysis of the optimal inventory strategy outlined in Scenario 8, Figures 6 through 8 present the inventory behavior over time for one representative SKU from each demand category—A, B, and C—managed by the Logistics Service Provider (LSP) in the EMEA region. These visualizations showcase inventory performance over 140 days, comparing actual inventory levels with the original safety stock targets and the lower bounds introduced in this study. The objective is to assess how effectively the proposed policy maintains service levels while minimizing holding costs for each demand segment.

In Figure 6, we examine SKU A1, which belongs to Category A and represents the highest demand volume among the nine SKUs. This SKU is managed under a continuous review policy with daily replenishment. Initially, inventory levels exhibit fluctuations, but they stabilize after Day 28, coinciding with the full accounting of lead time. Once the system is steady, inventory remains above the lower bound and closely aligns with the safety stock targets. This underscores that high-frequency replenishment effectively maintains availability for high-demand items while preventing excessive stock accumulation. The early fluctuations emphasize the importance of considering lead time when implementing new inventory strategies, while confirming that the model quickly adjusts to stabilize performance following the lead time.

Figure 6. Inventory Quantity Over Time at EMEA LSP for SKU A1 (Category A)



In contrast, SKUs from Categories B and C, which are managed using periodic review policies with intervals of six and four days, respectively, display a more cyclical inventory pattern. SKU B1, as shown in Figure 7, experiences regular peaks and troughs in line with its replenishment cycles while consistently remaining within the desired inventory range. Similarly, SKU C1, depicted in Figure 8, although characterized by smaller volumes, maintains inventory levels above the lower bound with minimal variability, demonstrating a strong alignment with the demand pattern. These findings validate the effectiveness of longer replenishment cycles for lower-demand products, allowing for reduced operational costs while maintaining service reliability. These insights emphasize that Scenario 8's tailored strategy provides a robust and adaptive inventory solution across various demand profiles.

Figure 7. Inventory Quantity Over Time at EMEA LSP for SKU B1 (Category B)

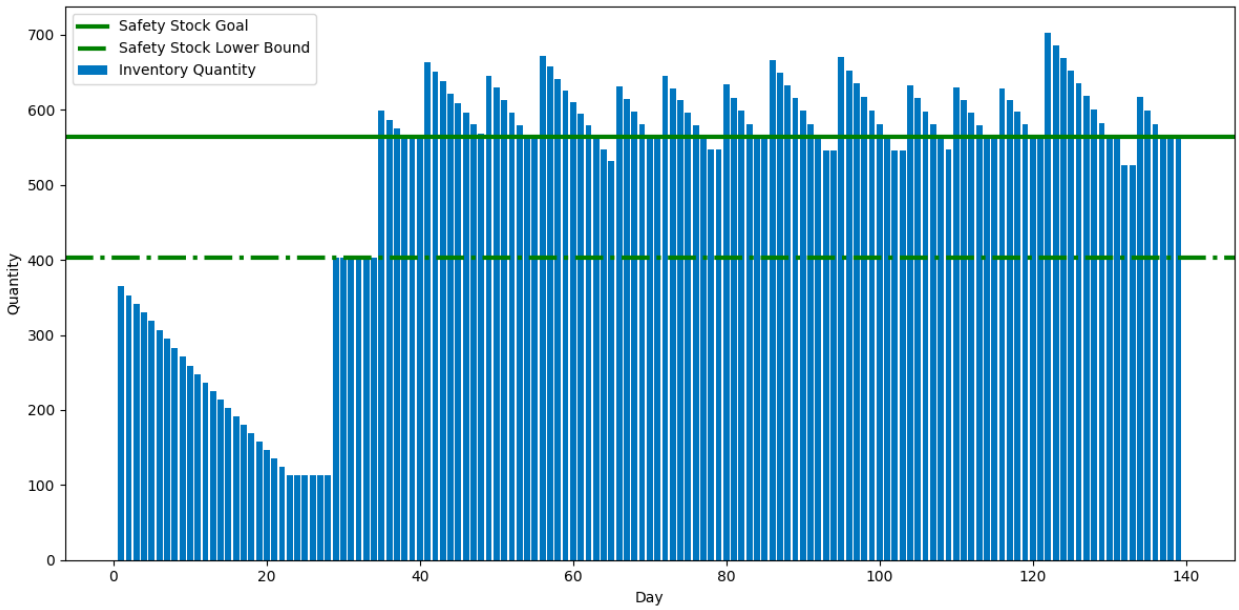
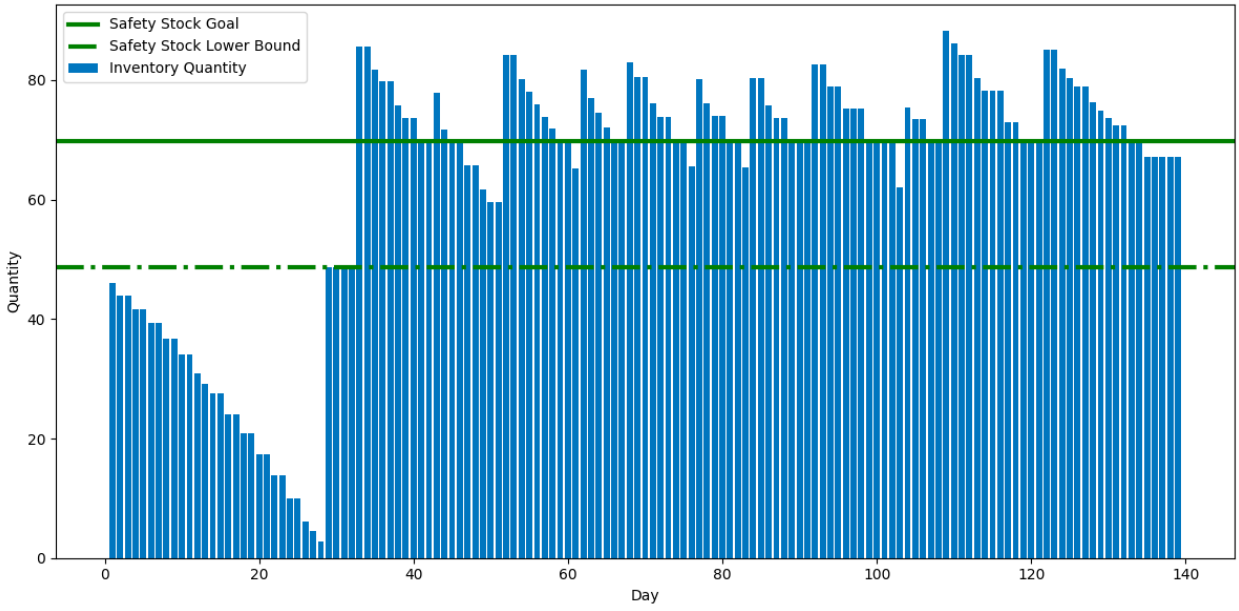


Figure 8. Inventory Quantity Over Time at EMEA LSP for SKU C1 (Category C)



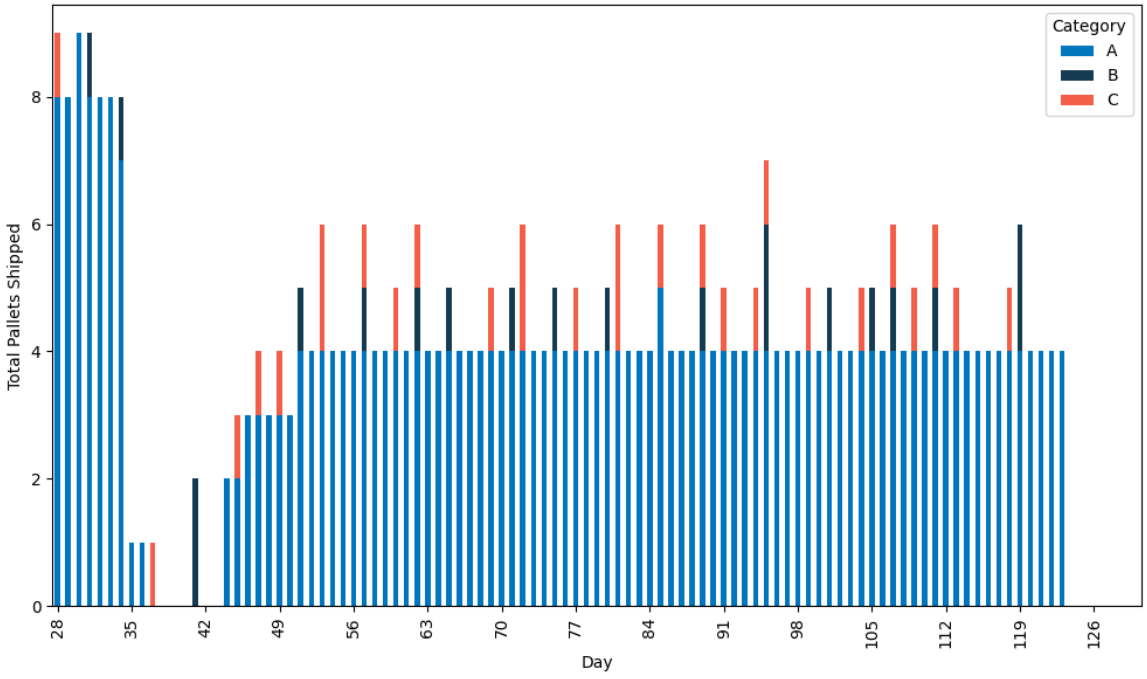
4.4. Proposed Replenishment Strategies

Building on the findings of Scenario 8—the optimal strategy derived from quantitative and qualitative assessments—this section articulates the proposed replenishment methodology. Within this framework, Category A SKUs, characterized by high demand, will be replenished on a continuous review basis (0-day cycle). In contrast, SKUs classified as Category B and Category C, representing medium and low demand, will be managed according to periodic review policies of 6 and 4 days. This configuration

balances responsiveness with operational efficiency by synchronizing replenishment frequency with the specific demand profiles of each SKU category.

Figure 9 illustrates the total number of pallets shipped daily from the Global Distribution and Replenishment Center (GDRC) to the EMEA Logistics Service Provider (LSP), categorized by product group. This chart visualizes outbound logistics activities over time, quantitatively capturing the effects of staggered replenishment cycles across the distinct SKU categories. Notably, the early segment of the time horizon reveals pronounced spikes corresponding to the initial fulfillment phase, subsequently giving way to a more stable and regular shipping pattern as the model attains a steady state following the consideration of lead time.

Figure 9. Daily Pallets Shipped from Global Distribution and Replenishment Center (GDRC) to the EMEA Logistics Service Provider (LSP) by category



The graph elucidates that Category A consistently drives daily shipment volumes in alignment with its 0-day replenishment cycle. Conversely, Categories B and C display shipment activity at defined intervals, consistent with their respective 6-day and 4-day cycles. This pattern is further supported by the replenishment frequency data in Table 8, which indicates that Category A SKUs were replenished on more than 80% of the days within the 96-day range, with some reaching nearly 90%. In contrast, Category B SKUs were replenished between 4% and 14.6% of the time, while Category C SKUs ranged from 4.2% to 20.8%. These results confirm that the replenishment cadence in Scenario 8 reflects demand dynamics and

results in a more predictable and consolidated shipping structure. Such consistency enhances planning for warehouse operations and transportation, supporting the objective of reduced labor burden and optimized logistics at the Global Distribution and Replenishment Center.

This systematic pattern supports objectives related to shipment consolidation and the optimization of warehouse workload. By eschewing daily replenishment for all SKUs, the proposed strategy effectively mitigates labor intensity and enhances container utilization, all while meeting requisite inventory and service levels. These insights substantiate that the replenishment strategy proposed in Scenario 8 provides a viable framework for enhancing operational efficiency and achieving cost control, particularly in multi-node, category-diverse supply chains such as that of Cordis.

Table 8. Number and Percentage of Days Replenished per SKU at EMEA LSP (Scenario 8)

SKU	*Replenishment Days*	*Days Replenished (%)*
A1	83	86.5%
A2	80	83.3%
A3	86	89.6%
B1	4	4.2%
B2	14	14.6%
B3	4	4.2%
C1	10	10.4%
C2	4	4.2%
C3	20	20.8%

As shown in Table 9, the SKUs in Category A, which are managed under a continuous review policy, had their Economic Order Quantity (EOQ) and reorder point (s) calculated using Equations 2 and 3 from Section 3.3, Inventory Review Policy. These calculations are essential for high-demand products, as they enable timely and efficient replenishment based on real-time inventory positions, reducing the risk of stockouts while minimizing holding and ordering costs. In contrast, Table 10 for SKUs in Categories B and C, which operate under periodic review policies, had their reorder points calculated using Equation 4, which accounts for both lead time and review period to determine the appropriate inventory target level at each review cycle.

Table 9. Category A SKUs Continuous Review Policy Parameters

SKU	*Replenishment Cycle*	*EOQ*	*s*
A1	0	241	8,274
A2	0	312	5,086
A3	0	305	5,123

Table 10. Categories B and C SKUs Periodic Review Policy Parameters

SKU	*Replenishment Cycle*	*s*
B1	6	161
B2	6	329
B3	6	40
C1	4	35
C2	4	26
C3	4	82

This distinction is crucial, as it ensures that the replenishment approach aligns with each SKU's demand profile and operational role. High-frequency, demand-driven policies for Category A maximize service levels for critical items, while the structured periodic policies for Categories B and C support order consolidation and cost control. Applying the appropriate mathematical models to each review policy improves overall supply chain responsiveness, efficiency, and financial performance, reinforcing the strategic foundation of Scenario 8.

4.5. Discussion

The findings of this study underscore the value of differentiated inventory strategies in the medical device industry, where high service levels and regulatory precision must coexist with cost efficiency. Cordis, like many companies in this sector, faces the challenge of balancing rapid fulfillment for life-saving products with the need to control growing operational costs. By applying a formula-based safety stock approach and aligning replenishment cycles with SKU demand profiles, this study offers a path forward—one that supports Cordis's Breakthrough Lean (BTL) goals while preserving the service standards required in healthcare supply chains.

Our analysis revealed that while continuous review policies may minimize total costs in some cases, they also introduce significant operational strain, particularly for slower-moving products. Scenario

8 emerged as the most practical and cost-effective solution, offering a 35% reduction in operational costs through a hybrid replenishment policy. This result highlights a key insight for the broader industry: standardizing inventory policies across all products may lead to avoidable inefficiencies, especially in diverse portfolios like those in medical device manufacturing.

Moreover, these findings are particularly relevant in an era where agility, visibility, and cost control are more important than ever. As global supply chains continue to face disruptions, the ability to dynamically adjust safety stock and replenishment policies based on real-time demand and risk is becoming a competitive necessity, not just a cost-saving tactic. By adopting the strategies proposed in this capstone, Cordis can lead the way in building a smarter, more resilient supply chain, setting a benchmark for other firms operating in similarly high-stakes, regulated industries.

4.6. Limitations

While the findings of this project provide actionable insights for improving Cordis's inventory and replenishment strategy, several limitations should be acknowledged. One key limitation involves the treatment of penalty costs. The sponsor company had not previously established specific cost values for backorders or for failing to meet safety stock targets. To address this, we assumed assigning each penalty a value equivalent to the product's price. While this approach allowed the model to prioritize service level compliance and avoid stockouts, it may not accurately reflect the true business impact of backorders or inventory shortages.

Another limitation is related to data availability and scope. The historical data provided for this analysis covered only a single quarter of operations, which may not fully capture seasonal demand fluctuations, promotional effects, or variability associated with longer time horizons. As a result, safety stock levels and replenishment recommendations might differ if full-year or multi-year data were available. Additionally, the average daily demand for some SKUs within Categories B and C showed limited differentiation, which could affect the accuracy of ABC classification and lead to less distinct policy outcomes.

Finally, modeling constraints also introduced simplifications. For instance, applying a fixed lower bound to safety stock across all scenarios may have limited the model's flexibility in optimizing replenishment frequency and cost. Removing or adjusting this bound could yield different trade-offs between operational efficiency and service performance. As such, future work could explore sensitivity analyses on penalty values, incorporate seasonality, and further refine SKU categorization to validate and enhance the robustness of the proposed strategy.

4.7. Managerial Insights

The scenario-based optimization undertaken in this project provides Cordis with a clear, data-driven roadmap to enhance inventory and replenishment efficiency. The most cost-effective configuration, Scenario 8, revealed that implementing a hybrid inventory policy delivers substantial benefits. By utilizing a continuous review policy for high-demand SKUs and periodic replenishment cycles of six and four days for medium- and low-demand SKUs, respectively, Cordis achieved a remarkable 35% reduction in operational costs compared to the current baseline. Importantly, this approach maintained inventory levels within acceptable service thresholds across all SKUs. By aligning replenishment frequency with demand profiles, Cordis can reduce excess handling, consolidate shipments, and optimize container utilization without sacrificing service standards.

For managers, these outcomes underscore the necessity of transitioning from a uniform, one-size-fits-all inventory strategy to a segmented approach. Slower-moving items particularly benefit from decreased ordering frequency and shipment consolidation, which helps reduce labor efforts, lower transportation costs, and minimize inventory holding expenses. Conversely, high-demand SKUs require more responsiveness, which is achieved through frequent review cycles. This balanced approach ensures operational efficiency while upholding Cordis's commitment to high product availability in the medical supply chain.

To implement these findings, Cordis should consider operationalizing this hybrid strategy in phases, starting with a pilot program at the EMEA LSP and gradually expanding it across the global network. The optimization model created in this study can act as a decision-support tool to fine-tune replenishment policies based on region, product group, or service level targets. While some assumptions, such as fixed lead times or static safety stock targets, may limit the model's generalizability, the framework lays a scalable foundation for more adaptive and intelligent supply chain planning.

5. CONCLUSION

This capstone project focused on developing an inventory management and replenishment strategy to tackle Cordis's operational challenges, including high ordering frequency, generalized inventory levels, and rising supply chain costs. By categorizing SKUs into high, medium, and low-demand groups, we introduced a data-driven approach that combines differentiated replenishment policies with a formula-based safety stock strategy. Through an optimization model, we evaluated eight replenishment scenarios and found that a hybrid strategy, utilizing continuous review for high-demand SKUs along with periodic replenishment cycles of six and four days for medium and low-demand SKUs, respectively,

significantly enhances efficiency. The optimal configuration, Scenario 8, achieved a remarkable 35% reduction in operational costs compared to the existing policy.

A key managerial insight is the importance of aligning replenishment cadence with demand dynamics. Implementing periodic policies for slower-moving products facilitates shipment consolidation, reduces material handling, and enhances freight utilization, all without compromising service levels. This approach provides Cordis with a scalable pathway to decrease labor intensity, control ordering costs, and ensure product availability. Additionally, our findings advocate for a dynamic safety stock model that adjusts inventory targets based on variability and lead time, moving beyond the constraints of a static five-weeks-on-hand strategy.

While the results are encouraging, certain assumptions, such as fixed lead times and penalty cost estimates, could impact outcomes and warrant further validation. The analysis was also based on a single quarter of operational data, which may not fully reflect seasonal trends. Future work could benefit from incorporating rolling demand forecasts, testing the model with a broader SKU range, and exploring integration with real-time supply chain visibility tools.

The implications of this study extend beyond Cordis. Inventory management in the healthcare and medical device sectors is known to be complex and costly, particularly within consignment and just-in-time delivery frameworks. Inefficient inventory strategies contribute to escalating costs, fulfillment delays, and supply chain inefficiencies that can ultimately impact patient care. As hospitals face increasing cost pressures, suppliers must adopt smarter, leaner inventory systems. The framework developed in this project provides a practical, evidence-based solution that can serve as a model not only for Cordis but also for other medical manufacturers and healthcare supply chains seeking to reduce costs and enhance service performance in a demanding, high-stakes environment.

5.1 Future Research

Future research could explore the scalability of the proposed methodology by applying it to a larger set of SKUs across Cordis's full product portfolio. Expanding beyond the nine representative items used in this study would help evaluate whether the hybrid replenishment strategy and formula-based safety stock model remain effective under more complex and varied demand conditions. Additionally, extending the analysis to include multiple global distribution points—rather than focusing solely on the EMEA LSP—could reveal opportunities for network-wide consolidation, where synchronized replenishment policies across regions may enhance cost efficiency, improve service levels, and reduce backorders at a system-wide level.

Another valuable direction would be to integrate product shelf life or expiration constraints into the optimization model. This addition is particularly relevant for medical devices or consumables with time-sensitive usability. Accounting for perishability would allow for more realistic inventory planning, ensuring turnover is aligned with regulatory and clinical requirements. Future studies could also incorporate dynamic demand forecasting, SKU substitution logic, or real-time data integration to support more adaptive and intelligent inventory management strategies.

REFERENCES

Anjani, A., & Nizar, A. (2021). Inventory management and cost efficiency. *International Journal of Research in Business and Social Science (2147-4478)*, 10(2), 217–227. https://doi.org/10.20525/ijrbs.v10i2.1042

Correll, D. (2022, October 17). MIT SCM.260—Logistics Systems.

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

Global Healthcare Exchange. (2011). *Strategies for healthcare supply chain collaboration*. https://www.ghx.com/media/4odd4jrs/strategies_for_supply_chain_collaboration_1011.pdf

Hafnika, F., Farmaciawaty, D. A., Adhiutama, A., & Basri, M. H. (2016). Improvement of inventory control using continuous review policy in a local hospital in Bandung City, Indonesia. *The Asian Journal of Technology Management (AJTM)*, 9(2), 109–119. https://doi.org/10.12695/ajtm.2016.9.2.5

Hillstrom, N. L., & Malabanan, R. A. (2013). *Strategic inventory management of externally sourced medical devices* [Master's thesis, Massachusetts Institute of Technology]. DSpace@MIT. [Ihttps://dspace.mit.edu/handle/1721.1/81099?show=full]

Jonsson, P., & Mattsson, S. A. (2019). An inherent differentiation and system level assessment approach to inventory management: A safety stock method comparison. *International Journal of Logistics Management*, 30(2), 663–680. https://doi.org/10.1108/IJLM-12-2017-0329

Setyaningsih, S., & Basri, M. H. (2013). Innovations in logistics management. Retrieved from https://ijimt.org

Tjhin, P. C., & Pandey, R. (2006). Inventory management strategy for the supply chain of a medical device company [Master's thesis, Massachusetts Institute of Technology]. DSpace@MIT. [https://dspace.mit.edu/handle/1721.1/77474]