Machine Learning in Prescribing Optimal Target Inventory Levels for FMCG

by

Wenchao Zhang

Bachelor of Science in Maritime Studies, Nanyang Technological University, 2015

and

Firanza Fadilla

Bachelor of Science in Chemical Engineering, Institut Teknologi Bandung, 2017

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 2024

Signature of Author: ______________________________

Wenchao Zhang
Department of Supply Chain Management
May 10, 2024

Signature of Author: ______________________________

Firanza Fadilla
Department of Supply Chain Management
May 10, 2024

Certified by: ______________________________

Dr. Ilya Jackson
Postdoctoral Associate
Capstone Advisor

Certified by: ______________________________

Dr. Devadrita Nair
Postdoctoral Associate
Capstone Co-Advisor

Accepted by: ______________________________

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

Machine Learning in Prescribing Optimal Target Inventory Levels for FMCG

by

Wenchao Zhang

and

Firanza Fadilla

Submitted to the Program in Supply Chain Management

on May 10, 2024 in Partial Fulfillment of the

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

ABSTRACT

For our project sponsor, a leading global Fast Moving Consumer Goods (FMCG) company, the challenge lies in optimizing inventory management strategies to balance avoiding stockouts with minimizing costs associated with excess inventory. This capstone project explores the use of machine learning to optimize target inventory levels, thus improving service levels. Developing features and creating labels using company data, we examined various machine learning techniques such as Random Forest and Extreme Gradient Boosting or XGBoost, and fine-tuned hyperparameters to develop the most accurate binary classification model. The Random Forest model predicts stockout events for each stock keeping unit (SKU) and warehouse combination with 91% accuracy, evaluated by f1-score and recall. We further conducted sensitivity analyses on the best performing model to discover the elasticity between service levels and various scenarios of increased or decreased target inventory levels. Based on our results, we provide the company recommendations to adjust safety stock calculations during seasonal transitions, reducing excessive inventory by 15% while improving the overall target service level across all SKUs to 97%.

ACKNOWLEDGEMENTS

I am immensely grateful to my capstone project partner, Firanza Fadilla, for his creative spirit and collaborative approach. His dedication was crucial to our project's success. My sincere thanks to Dr. Ilya Jackson and Dr. Devadrita Nair for their expert guidance and support in machine learning and data science. Their passion deeply influenced my aspirations to excel in data science and modeling. I also thank Dr. Sean Willems, my professor of Supply Chain Analytics, for his valuable insights into multi-echelon inventory optimization, enhancing my understanding of complex supply chain issues.

Wenchao Zhang (Philip)

I deeply appreciate our advisor and co-advisor, Dr. Ilya Jackson and Dr. Devadrita Nair, for their support, guidance, and commitment to keeping our project on track and within scope. I am also grateful to the sponsor company for granting us access to valuable, high-quality data, which provided a strong foundation for our research and allowed us to begin effectively. Special thanks to Pamela for her keen attention to detail and ongoing encouragement during the writing process. I am thankful to my capstone partner, Philip, for his crucial role in helping us start strong and for his creative input in developing our application.

Firanza Fadilla

LIST OF TABLES

LIST OF FIGURES

TABLE OF CONTENTS

1. Introduction

In the dynamic environment of the North American FMCG industry, where demand patterns are in constant flux, the importance of precise inventory management cannot be overstated. Several factors contribute to this demand variability. Seasonal shifts, often tied to promotional seasons, can lead to heightened demand at specific times of the year. The industry also feels the ripples of potent advertising campaigns or the introduction of new products, causing short-term demand surges. Economic dynamics, whether downturns or upswings, directly shape consumer purchasing power and, by extension, the demand for FMCG products. Moreover, contemporary trends, societal norms, and values, often influenced by social media and celebrity endorsements, have a significant bearing on the popularity of certain products.

For a global leader like our project sponsor, a leading global CPG company, the challenge is not merely to meet this multifaceted demand but to do so in an optimized manner. The company's inventory management strategies, while efficient in their current form, present an opportunity for further optimization. Setting the target inventory level appropriately is critical; if set too low, the company risks stockouts, failing to meet customer demands and service levels. On the other hand, setting it too high escalates inventory holding costs, which include more than just the expenses of storage. Excessive inventory can lead to tied-up capital, risk of product obsolescence, and depreciation, highlighting the importance of a balanced approach. This delicate trade-off between maintaining sufficient inventory to meet demand and minimizing holding costs creates a significant opportunity for optimization. Our project aims to refine its approach to determining the optimal target inventory level, employing advanced analytical methods to better align with the dynamic nature of market demands and operational realities, thereby enhancing overall efficiency and reducing unnecessary financial burdens.

The integration of big data analytics into inventory control is a pivotal upgrade for the project sponsor's existing methods. This approach significantly enhances forecasting accuracy by analyzing large volumes of data on market trends, customer behaviors, and seasonal patterns, thus improving decision-making regarding inventory levels and production schedules. Additionally, big data analytics aids in optimizing the supply chain by providing insights into supplier performance and operational effectiveness, incorporating external market data and real-time feeds like social media sentiment, thereby offering a more dynamic and responsive inventory management system.

The project's central focus is on improving upon the current statistical model, which combines demand forecasting and simulation, by employing Machine Learning (ML) models to prescribe more accurate target inventory levels. This involves leveraging multivariate data, both endogenous and exogenous, to refine the dynamic inventory control model. The following research questions stem from this problem:

- How can the current target inventory level setting be optimized at a granular level for each SKU and site?
- What key variables or features should be incorporated into a multivariate dynamic inventory control model?
- Which machine learning model can prescribe a more accurate target inventory level than the current model, and how?

The primary objective of this project was to develop a comprehensive ML-based model that can accurately prescribe target inventory levels, considering a range of influencing factors. This involved:

- **Developing an advanced ML classification model**: We constructed an advanced classification model that integrates a broad spectrum of internal and external variables into the organization. The goal was for this model to surpass the performance of existing forecasting and simulation models in terms of accuracy and reliability.
- **Conducting Sensitivity Analyses**: We performed sensitivity analyses using the trained classification model to explore the relationship between service levels and safety stock quantities. This analysis helped in understanding the impact of various factors on inventory requirements and facilitated more informed decision-making.
- **Performing a Root Cause Analysis**: We conducted a deep dive into specific SKUs and sites that presented opportunities for improvement. This analysis focused on identifying and understanding the underlying causes of operational inefficiencies, providing a clear path for targeted interventions.

Our research findings led to significant managerial insights and actionable outcomes for the sponsoring company. Key findings include:

- **Reduction of Excessive Safety Stocks**: We identified approximately 2,180 units of excessive safety stock across 27 SKUs at various facility sites. This discovery allowed for the optimization of inventory levels and the potential reduction of associated costs.
- **Repositioning Safety Stocks to Enhance Service Levels**: By strategically repositioning 6,155 units of safety stock across 24 facility sites, we saw an increase in the current service level from 95.2% to 96.7%.
- **Addressing the Landslide Effect**: Our analysis revealed a "landslide effect" associated with the forward coverage inventory policies during seasonal transitions of consumer-packaged goods (CPG). We proposed a novel method for setting target inventory levels that minimizes out-of-stock (OOS) situations when transitioning from high to low season and reduces safety stock costs when moving from low to high season. The implementation of the optimized target inventory levels is expected to improve service levels by an additional 6-7% and reduce inventory costs by 20%.

2. State of The Practice

Managing inventory level for FMCG products, presents unique challenges due to a combination of factors: demand variability, inventory management constraints, and consumer behavior. To achieve the primary objective of this capstone project — generating reliable demand forecasts that can be utilized to identify target inventory levels using ML techniques — we undertook a comprehensive literature review concentrated on three key areas:

1. FMCG Products Market Dynamics
2. Inventory Management in FMCG Industry
3. Machine Learning Techniques

2.1 FMCG Products Market Dynamics

The global FMCG market is experiencing significant growth, with the Asia-Pacific region leading in demand. In contrast, North American markets have seen stagnant growth due to market maturity. Economic factors are also influencing the market; stable disposable incomes in developed countries and increased purchasing power in developing nations have made FMCG products more affordable globally (Technavio, 2023).

The FMCG market is dynamic and competitive, grappling with challenges surrounding cost, convenience, and sustainability. High production and distribution costs compel manufacturers to balance affordability and profitability while facing competition from sustainable alternatives, which have gained popularity due to environmental concerns. Moreover, e-commerce's emergence as a preferred shopping platform boosts FMCG sales online. Looking ahead, the market is poised for further growth, driven by demand for convenient, eco-friendly products and the expansion of e-commerce (Statista, 2023).

However, challenges persist, including competition from other products and operational costs. To stay competitive, companies must innovate and emphasize sustainability in their products and marketing approaches, aligning with evolving consumer preferences (Statista, 2023). The market also faced challenges in 2020 due to COVID-19 lockdowns, leading to manufacturing disruptions and distribution channel breakdowns. The introduction of vaccines and lifting of lockdowns in Q3 2020 has led to a market recovery, with resumed manufacturing and adherence to new social distancing norms. The market is expected to grow as the focus on hygiene remains high (Technavio, 2023).

2.2 Inventory Management in FMCG Industry

Fast-moving consumer goods (FMCGs) encompass everyday items that are rapidly sold at a low cost to a broad consumer base (Lee & Siddiqui, 2023). To address the rapidly changing demand and manage the challenge of quick product turnover, companies need to keep ideal inventory levels, both in-store and at various stages of their supply chains (ITC Infotech, 2020). It is particularly crucial for FMCG companies operating on a built-to-stock model, where they need to allocate resources to finished goods based on projected demand, aiming to limit their financial risk. Those built-to-stock firms experiencing marked seasonal demand are required to initiate their planning and inventory accumulation sooner and often with less data compared to their counterparts in the non-seasonal CPG sector (Gundogdu & Maloney, 2019).

A key challenge in managing the supply chain of FMCG is identifying the ideal inventory levels at each stage of the supply chain. This is crucial to prevent stockouts, which may result in losing customers, and to avoid surplus inventory, which can lead to elevated storage and financial expenses (Inderfurth, 1991). The FMCG industry's high-volume, low-margin nature demands efficient inventory management to maintain uninterrupted supply and meet customer expectations. Seasonal variations and market fluctuations further complicate forecasting, making it crucial to monitor trends and historical data for informed inventory level adjustments (Rapidor, 2023).

2.2.1 Performance Metrics for Inventory Control

Global supply chains are vulnerable to disruptions that adversely affect their efficiency. Typically, these disturbances result in failure modes that compromise the supply chain's capacity to deliver goods and services as scheduled. Consequently, businesses involved in various supply chains are eager to develop resilience against such disruptions and their associated failure modes to ensure on-time delivery and maintain competitiveness (Carvalho et al., 2022). To establish an effective supply chain, businesses often track a variety of supply chain metrics or key performance indicators (KPIs). These KPIs are chosen for their significance and influence on the efficiency of the company's supply chain processes. The specific metrics differ across various levels and functions within the supply chain (Lohman et al., 2004).

For inventory control, service levels constitute arguably the most important performance measures. The fill rate in particular (which determines the percentage of demand satisfied directly from stock-on-hand) is the most commonly used measure in industry, as it translates directly to the customer service level achieved (Teunter et al., 2017).

Product availability indicates a company's capacity to fulfill customer orders from its existing inventory. If a customer places an order when the product is unavailable, this leads to a stockout. There are multiple methods to assess product availability (Chopra & Meindl, 2016):

- The product fill rate (FR) represents the proportion of product demand met using inventory stock. It is essentially the likelihood of meeting product demand from the available inventory. The fill rate is better assessed over specified amounts of demand rather than over time.
- Order fill rate refers to the percentage of orders that are satisfied using the available inventory. The order fill rate should be evaluated based on a specific number of orders, rather than over a period of time. In situations involving multiple products, an order is considered filled from inventory only if all items in the order are available in the inventory.
- Cycle Service Level (CSL) is the proportion of replenishment cycles that end with all customer demands fulfilled. A replenishment cycle refers to the time span between two consecutive replenishment deliveries. CSL corresponds to the likelihood of avoiding stockouts during a replenishment cycle. It should be calculated over a specified number of these cycles.

In this capstone, the sponsoring company uses CSL as a key metric and a benchmark to measure the performance of customer delivery. This approach, known as Service Level-Driven Planning (SLDP), assists organizations in identifying and reaching a consensus on their optimal inventory level, or the "sweet spot". This spot is where the precise amount of product or material is designated to meet the desired service levels throughout the replenishment lead time. SLDP offers a concrete figure that forecasts the probable service level for each stocked item and for the overall inventory (Hartunian, 2015).

2.2.2 Inventory Models

Effective inventory management involves the implementation of models, techniques, and methods that facilitate proper resource utilization, control costs in operations, and aid in making informed decisions. These models encompass a range of policies and controls that track inventory levels, ascertain the optimal levels to be sustained, the appropriate timing for restocking inventory, and the ideal order sizes. The literature identifies two primary categories of these models based on the nature of demand: (i) mathematical deterministic model (regular demand) and (ii) stochastic model (probabilistic demand) (Vidal, 2023).

The transition to policies like periodic review (R, s, S) marked significant advancements in handling stochastic single-product inventory control systems. This approach merges the (s, S) and (R,S) inventory systems. The concept involves inspecting the inventory levels at every interval of R time units. Should the inventory level be at or fall below the reorder threshold s, an order is placed to increase the level up to S. Conversely, if the inventory level is above s, no action is taken until the subsequent review period arrives (Silver et al., 2017). The value of S is determined based on service measure, which for this research is the service level.

Demand forecasting plays a key role in inventory management (Gardner, 1990). Moreover, it is key in determining the order-up-to-level (S) under a periodic review inventory control system. Order-up-to inventory level needs to ensure that the inventory levels can cover the average demand of the product and buffer against any uncertain demand fluctuation in between the review period (R). Therefore, S has two key components, average demand over the review period and the safety stock, also known as buffer stock. The forecast is used to calculate the safety stock component of the order-up-to-level. To do so, the demand for a product is forecasted using the historical demand of the product. Once the forecast is obtained, the root mean squared error (RMSE) of the forecast is used to assess if the forecast explains the variability in demand (Hyndman & Koehler, 2006).

The lower the RMSE, the better it explains the demand variability. The RMSE is used to calculate the safety stock component of order-up-to-level in the periodic review control system. These models, however, often operate under assumptions that may not align with real-world complexities in terms of (1) assuming demand distribution (2) assuming IID (independent and identically distributed) demand. Demand forecasting serves as a foundational tool for setting target inventory levels, employing a variety of methods and techniques. This approach ensures inventory optimization by accurately adjusting stock levels accordingly.

2.3 Target Inventory Level Determination: A Review

When determining the target inventory level, the methods can broadly be categorized into two main types: the classic model and the data model, as illustrated in Figure 1.

2.3.1 Classic Model for Optimal Target Inventory Level Setting

The classic method of determining the optimal target inventory level involves a process that starts with demand forecasting, then hypothesizes a particular demand distribution, follows with simulations, and concludes with establishing a target inventory level (Seyedan et al., 2023). This method, though fundamental, is limited by its presumptions about demand, often envisaging it as deterministic or adhering to a preconceived probability distribution, a notion that may be unsuitable in real-world scenarios where demand patterns are unpredictable or ambiguous (Clausen & Li, 2022).

Figure 1

Target Inventory Level Literature Structure

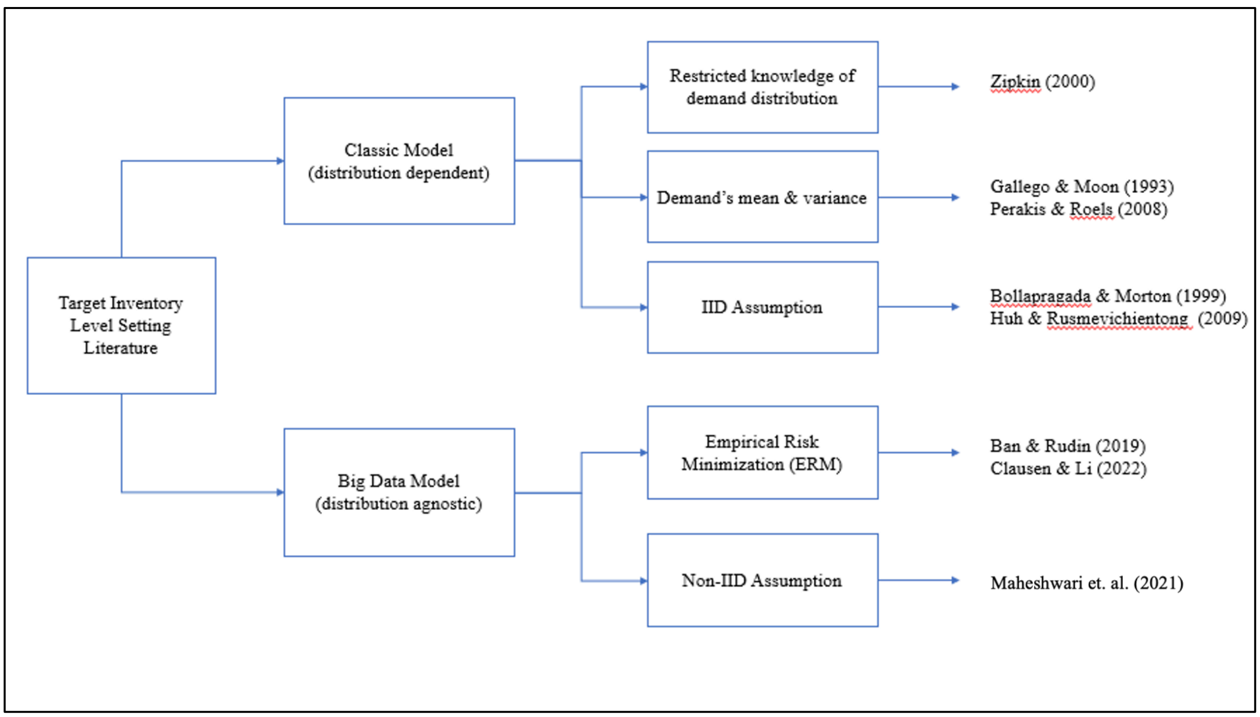

Pioneering work in this domain, such as that of Zipkin (2000), includes models that operate with restricted knowledge of demand distribution. For instance, Gallego and Moon (1993) developed newsvendor models utilizing only demand's mean and variance, refining the Scarf ordering rule. Perakis and Roels (2008) also highlighted the significance of distributional moments, like mean and variance, in their exploration of the newsvendor model. Zhang & Yang (2016) investigated multi-item newsvendor models employing historical univariate demand data.

2.3.2 Limitation of the Classical Model in Handling IID (independent and identically distributed) Demand Data

The classical model, exemplified by the newsvendor model, is instrumental in fostering new methodological insights. However, its effectiveness is predominantly confined to scenarios involving independent and identically distributed (IID) demand data. This limitation becomes apparent when considering the advancements in data-driven inventory models and solution algorithms, particularly within the realm of dynamic inventory models (Clausen & Li, 2022).

For instance, Bollapragada & Morton (1999) developed a heuristic specifically for the (s, S) inventory model that effectively addresses IID demand data scenarios. Similarly, Huh & Rusmevichientong (2009) drew inspiration from the newsvendor model to create a data-driven multiple period order-up-to level inventory control model. This model aims to optimize total costs and profits over a planning horizon, considering factors like order cost, holding cost, and lost sales cost, and is particularly adept when applied to IID data, as demonstrated by their univariate algorithm that achieves asymptotic optimality in such contexts.

These developments underscore a critical limitation of the classical model: its reliance on the assumption of IID demand data, which restricts its applicability and adaptability in more complex, real-world scenarios where demand patterns may not conform to the IID criterion. In this paper, the classic model can serve as a foundational baseline for comparison with big data models in determining the optimal target inventory level.

2.3.3 Big Data Approach in Inventory Management

In contrast to classic models, the big data model does not assume any specific distribution of demand data. It leverages machine learning techniques to uncover the relations between demand and various observed features, demonstrating robust performance even with non-IID demand data (Maheshwari et al., 2021). As a relatively nascent area of research, the big data-driven inventory model has seen limited but insightful studies.

Ban & Rudin (2019) made a significant contribution by applying this model to situations where businesses possess extensive historical demand data and a broad range of demand-influencing features, utilizing distribution-agnostic machine learning algorithms. This method proficiently copes with data-intensive scenarios and delineates clear performance metrics for out-of-sample costs. Further advancing this field, developed a Neural Network-based formulation of the big data-driven newsvendor model, utilizing the Empirical Risk Minimization (ERM) principle. Building upon this, Clausen & Li (2022) extended this research by employing ERM to craft a dynamic, big data-driven order-up-to level inventory model, complete with a custom-designed machine learning algorithm. Their results underscore the efficacy of this integrated approach, showing potential cost reductions of up to 60% over the most effective univariate models and up to 6.37% over the top big data-driven benchmark models.

This progression in the big data model's development highlights its capacity to handle complex demand data scenarios more effectively than traditional models, marking a significant advancement in inventory management research. As our sponsoring company demonstrates non-IID demand data, we will delve deeply into this model as a preferred technique to determine the target inventory level.

Figure 2

Different Types of Machine Learning (Sanjeevi et al., 2017)

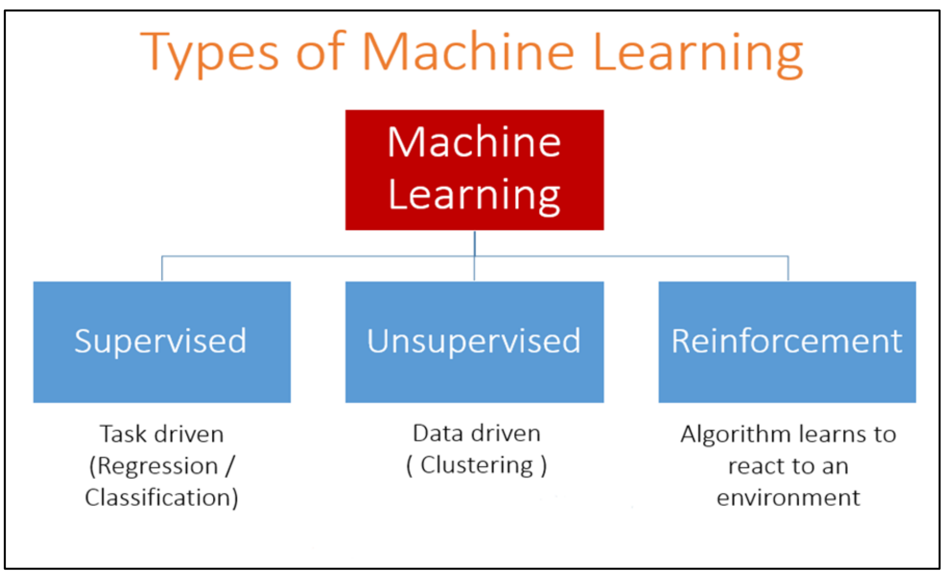

2.4 Applying Machine Learning to The Supply Chain

Machine learning, a subset of artificial intelligence, involves developing algorithms that learn from data to solve various problems. These techniques derive insights from historical data to

address similar challenges in the future (Michalski et al., 1983). Priore et al. (2019) identified key machine learning techniques for inventory policies, including artificial neural networks (ANN), case-based reasoning, support vector machines (SVM), and reinforcement learning.
In the increasingly complex and dynamic field of supply chain management, machine learning has become a crucial tool, as highlighted by de Melo (2019). This review examines the application of various machine learning techniques in inventory management, a key aspect of supply chain operations. According to Sanjeevi et al. (2017), machine learning algorithms are generally categorized into three types: supervised learning, unsupervised learning, and reinforcement learning, as shown in Figure 2.

2.4.1 Supervised Learning

In supervised algorithms a training set with correct responses is provided and, based on this training set, the algorithm generalizes to respond correctly to all possible inputs (Marsland, 2014). Some of the common supervised ML methods are:

- **Decision Trees** use observations about certain actions and identify an optimal path for arriving at a desired outcome (de Melo, 2019).
- **Random Forest** operates as an ensemble learning method. It constructs multiple decision trees during training and generates predictions by averaging the outputs of these trees. Known for its versatility in handling diverse data types and distributions, Random Forest is robust and reduces overfitting risks. Deraz (2023) applied it to optimize inventory levels for fast-moving consumer goods (FMCGs), exploring its effectiveness in predicting weekly Economic Order Quantity (EOQ) under two scenarios: parallel (where data is certain and available) and sequential (where data is predictable). The study found that Boosted Decision Tree (BDT), another similar algorithm, yielded accurate results in both scenarios.
- **XGBoost (Extreme Gradient Boosting)**, although predominantly recognized for its classification capabilities, excels in regression tasks. It employs a gradient-boosting framework to enhance its predictive accuracy. By minimizing an objective function that amalgamates a convex loss function with a regularization term, XGBoost efficiently and accurately predicts continuous values (de Melo, 2019). Utilizing decision trees as its base models, this ensemble method integrates them for effective predictions. Gumus & Kiran

(2017) highlight its efficiency, while Namir et al. (2022) demonstrate its practical application in creating a decision support tool for dynamic inventory management.

- **Support Vector Machines (SVM)**, more commonly associated with classification tasks, are also adept at regression, termed Support Vector Regression (SVR) in this context. SVR employs SVM principles but focuses on fitting as many data instances as possible onto the regression line while limiting margin violations (de Melo, 2019). This approach is particularly effective in identifying intricate relationships within data sets. Carbonneau et al. (2008) have shown that SVM, along with recurrent neural networks, can provide highly accurate forecasts in real-world data sets, significantly improving inventory control.

2.4.2 Unsupervised Learning

In unsupervised learning, algorithms learn from unlabeled data, uncovering hidden patterns within datasets (de Melo, 2019). The **K-nearest neighbor algorithm**, for example, has been instrumental in improving supply chain production policies (Akhbari et. al, 2014). Furthermore, Neural Networks, especially **Artificial Neural Networks (ANNs)** such as autoencoders, have proven effective in models geared towards data-driven decision-making in supply chain environments. This efficacy is evident in the research conducted by Clausen & Li (2022)and Huber et al. (2019), demonstrating the significant impact of these learning algorithms in the supply chain sector.

2.4.3 Reinforcement Learning

Reinforcement learning is a data driven methodology that is gaining prominence in many economic research areas, particularly in inventory management. It involves an agent learning to make decisions within an environment to achieve specific goals (de Melo, 2019). Giannoccaro & Pontrandolfo (2002) provide an early introduction of reinforcement learning methods to inventory management. Since then, the field has seen significant theoretical advancements, with Kara & Dogan (2018), Perez et al. (2021), and Wang et al. (2012) applying newer reinforcement learning methods to tackle inventory management problems using univariate demand data.

This exploration of the relevant literatures has enabled us to pinpoint several approaches for determining the optimal inventory levels for the company. These include involving demand forecasts that incorporates exogenous variables and establishing target inventory levels based on this forecast, taking into account the constraints of inventory management policy.

3. Methodology

Following several discussions with our partner company, we agreed to focus on the central research question: how to utilize machine-learning-based demand forecasting to determine inventory levels to withstand various market changes such as seasonality. Our literature review led us to identify two methods for establishing the company's optimized inventory levels: developing a demand forecast that includes exogenous variables and determining target inventory levels based on this forecast and inventory management constraints. The project goal is for the optimized target inventory level to help our partner company lower its inventory and associated holding costs, while still maintaining the determined service levels. Our proposed methodology, in comparison to the company's current approach, is illustrated in Figure 3.

Figure 3

Proposed Methodology Compared to Existing Method to Generate Target Inventory Level

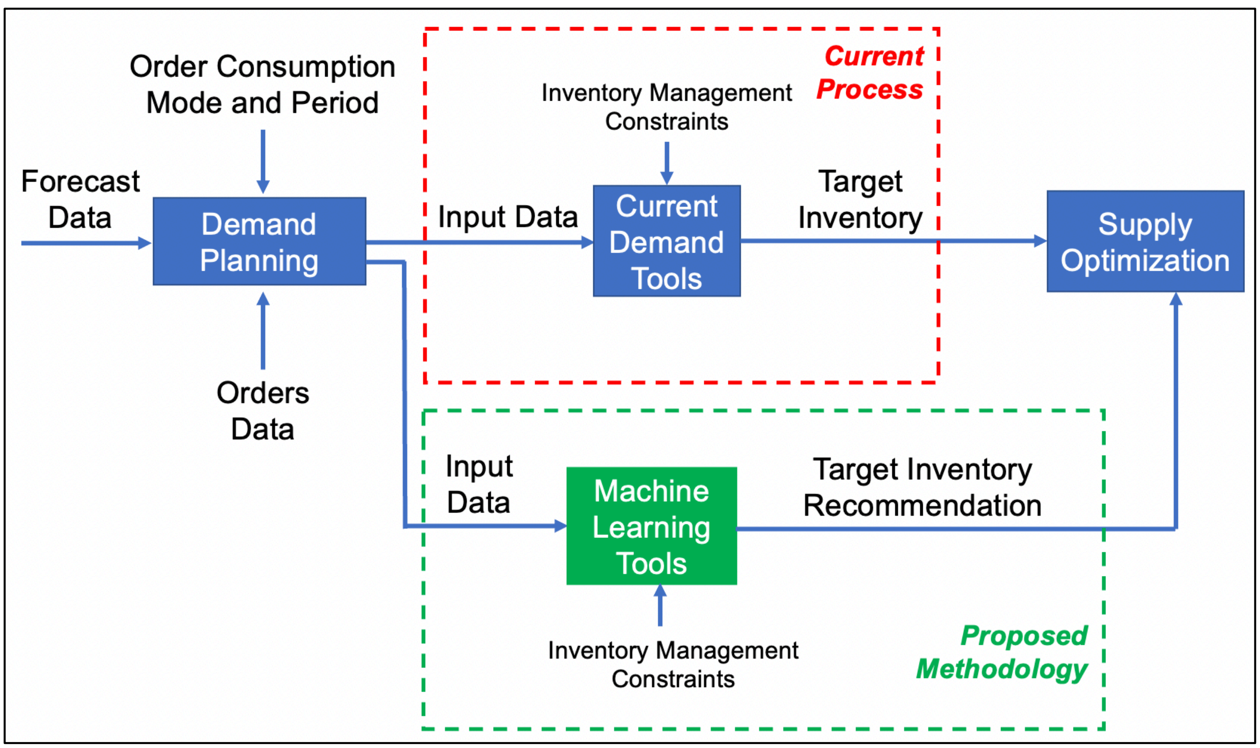

Our methodology is divided into 5 steps as shown in Figure 4.

Figure 4

Methodology Steps

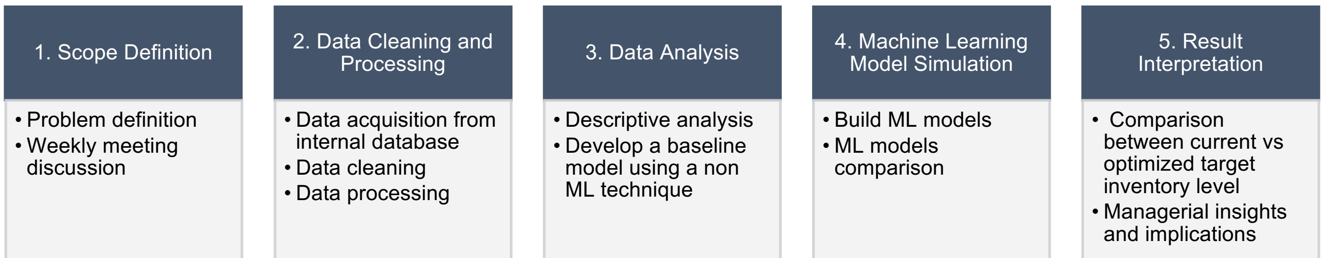

3.1. Scope Definition

Over several weekly meetings with the sponsoring company, we engaged in in-depth discussions focused on defining the problem at hand, thus ensuring a mutual understanding of the challenges and objectives. This also involved narrowing the scope of our project, making it more targeted and manageable. We also discussed the strategies for data access, considering the importance of data in driving our analysis and conclusions. Furthermore, we had extensive discussions about the methodology we proposed to ensure that our approach was both academically rigorous and relevant to the industry context.

3.2. Data Cleaning and Processing

This step included data acquisition, cleaning, and processing. For this project, The historical data and daily forecast were gathered from the company's internal database. These data took the form of inputs to the machine learning model, including historical open orders, historical shipments, historical target inventory level, and historical ending inventory. Data cleaning involved identifying and correcting (or removing) errors and inconsistencies from data to improve its quality, while data processing involved transforming raw data into a more usable and efficient format.

3.3. Data Analysis

In this process, we carried out a descriptive analysis and create an initial model utilizing a baseline technique. The descriptive analysis involved examining time series historical data, concentrating on metrics of central tendency, distribution, and identifying aspects like level, trend, and seasonality. Based on this historical data, we constructed a baseline model employing linear regression that incorporates exogenous variables.

3.4. Machine Learning Model Simulation

The main steps for this project were using historical data and exogenous variables as input to the multi-variate big data analytics model to generate a target inventory level as an output. We identified several machine learning techniques that could be applied to address this capstone's problem: XGBoost and Random Forest. For these different methods, we calculated the Forecast Value Added (FVA) for each technique compared to the baseline model. Based on the negative recall metric, we identified Random Forest as the best model.

3.5. Result Interpretation

In this step, we performed a comprehensive comparison of the current target inventory level against the optimized target inventory level. This comparison involved a detailed analysis of how the optimized levels differ from the existing ones and the potential impact of these differences on the business. We also extended the discussion to cover business and managerial insights, exploring the broader implications of these inventory levels on operational efficiency and overall business strategy.

4. Quantitative Analysis

In this chapter, we explore the quantitative aspects of our research. We started by presenting domain knowledge, followed by a descriptive analysis of data trends. Subsequent sections covered the selection and engineering of features, our data preprocessing approaches, and the selection and tuning of machine learning models suitable for inventory management. We also evaluated these models and present the classification results. Lastly, we analyzed the importance of various features using SHAP method, providing insights into their impact on model performance.

4.1 Exploratory Data Analysis

In our analysis, we focused on SKUs (Stock Keeping Units) within a certain size of the product line, encompassing over 500 SKUs across an 18-month period. This selection is aimed at capturing a wide spectrum of inventory dynamics and demand patterns, providing the robust dataset required for machine learning models. We simplified the company's multi-echelon supply chain into a two-echelon system consisting of three service warehouses (SW) and ten mixing centers (MC). SWs are multi-functional, managing both customer deliveries and replenishment to MCs, while MCs focus exclusively on customer service. This structure is significant given the differing lead times for replenishment (7 days) and customer shipments (20 days), as shown in Figure 5, with SWs directly fulfilling around 40% of customer deliveries, emphasizing their critical role in ensuring supply chain responsiveness.

Figure 5

Multi-echelon Supply Chain Network

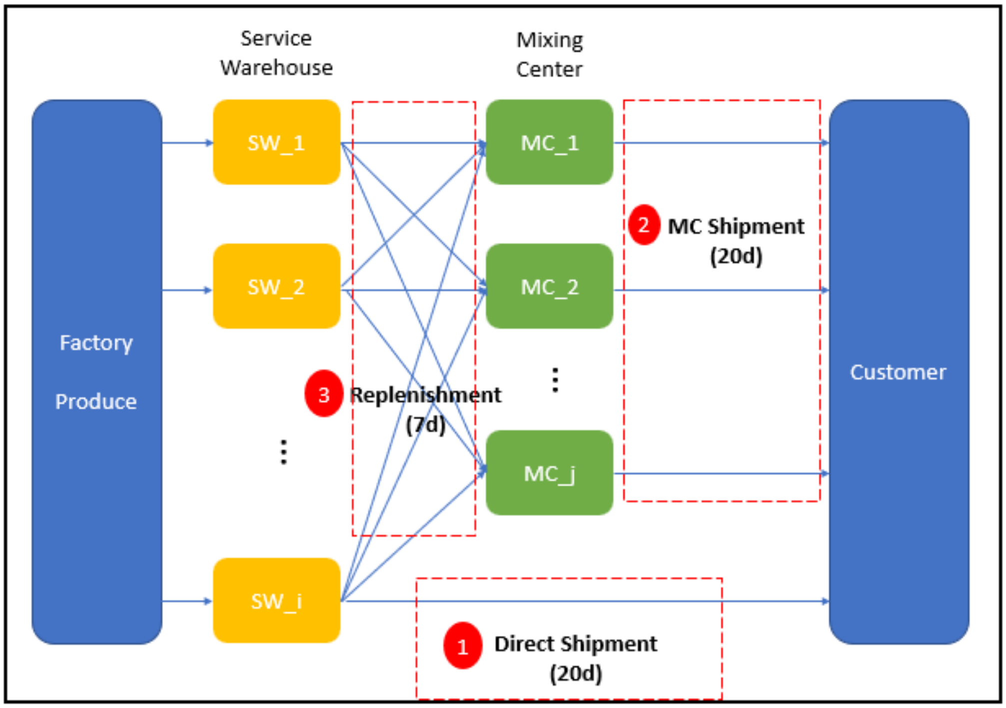

In line with the sponsored company's inventory management policy, we adopted a forward coverage approach. The company operates under a fixed forecast provided by the forecasting team, leaving the quantity of safety stock for each SKU at each site (SKU_site) and specific dates as the primary variable to manage demand variability. This safety stock aims to cushion against demand fluctuations over a predetermined forward coverage period, thereby ensuring optimal inventory levels and enhancing the robustness of the supply chain against unforeseen demand changes. This methodological foundation supported our objective of optimizing inventory levels to balance availability with cost efficiency, particularly within the context of a targeted review of forecast and demand data over an aggregated 20-day cycle. This period formed the basis for our safety stock calculation, aiming to cover the subsequent 20 days and establish a balanced inventory level that includes both safety and cycle stock considerations. Figure 6 presents the 20-day aggregated demand profiles for five selected high-volume SKUs and five selected low-volume SKUs over the past six months.

Figure 6

20-Day Aggregated Act Demand of 3 High Vol & 3 Low Vol SKUs Over 3 Months

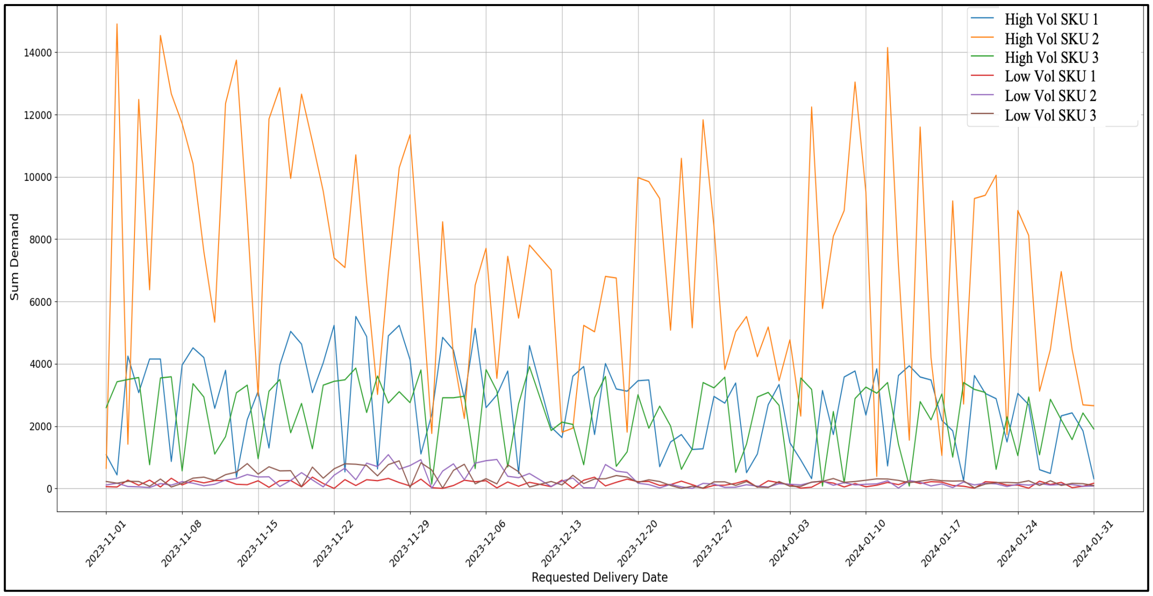

The graph reveals a significant disparity in the 20-day aggregated order quantities between the high and low-volume SKU-site combinations. We interpret this variability as an opportunity to ensure our model is both rigorous and representative by including a diverse range of SKU-site combinations.

4.2 Modeling

4.2.1 Feature Selection & Engineering

In preparing for the modeling phase, we selected a set of features from the company's internal database to predict inventory needs accurately.

- ***forecast_version_date:*** The date a future demand forecast is generated, with daily forecasts up to 60 days ahead and weekly thereafter.
- ***reqdlv_dt:*** The requested delivery date, indicating customer demand.
- ***sku_site:*** A categorical feature representing a unique combination of an SKU code and a site code, indicating the specific location where the product is stored.
- ***site_type:*** A categorical feature to differentiate whether the site code is a service warehouse ("SW") and a mixing center ("MC").

- ***safety_stock:*** A numerical feature representing the safety stock value at the requested delivery date for each sku_site
- ***ending_inventory:*** A numerical feature representing the actual inventory level at the requested delivery date for each sku_site
- ***sum_forecast:*** A numerical feature representing the aggregated forecasted order quantity over the next 20 days (customer shipment) or next 7 days (replenishment) at each sku_site

To enhance the model accuracy, we also included time series features below.

- ***Lag20:*** A time series feature that introduces the lag of sum_demand. A minimum lag of 20 days is selected based on the company's basis to capture sum_demand in determining safety stock.
- ***Month:*** Categorical variables representing the month of reqdlv_dt.
- ***Day of the Week and week of the year:*** Categorical variables representing the day of the week and week of the year of reqdlv_dt.
- ***is_quarter_end:*** A binary variable indicating the last day of a quarter.

After introducing all the features into our model, we determined the target variable that we aimed to predict. In this project, our goal is to predict the probability of stockouts using the current inventory levels and inventory management strategies. We defined the target variable as ***binary_label***, which assumes a value of 0 if the sum of the forecast value and safety stock is lower than the sum_demand, indicating a stockout situation where the safety stock is insufficient to meet the incoming demand. Conversely, it assumes a value of 1 when the sum of the forecast value and safety stock is greater than or equal to sum_demand, reflecting a scenario where the inventory is adequately protected against potential stockouts.

Having established the selection of features and the target variable, we moved on to the data preprocessing stage before feeding the data into the model. This step was crucial for ensuring that the imbalanced data was properly sampled and normalized for modeling, thereby enhancing the model's performance and accuracy.

4.2.2 Data Preprocessing

In the data pre-processing stage, we developed a pipeline to process both numerical and categorical data, essential for the model to accurately analyze the dataset's complexity. Given the dataset's imbalanced nature, we implemented sampling strategies to ensure a balanced class distribution: SMOTE (Synthetic Minority Oversampling Technique) for oversampling the minority class, and random under-sampling to reduce the prevalence of the majority class. These techniques helped achieve a more equitable representation of classes, thereby improving model accuracy and generalizability.

For time series data, we split the dataset based on a predetermined cutoff date (2023-11-30), dividing it into training and testing sets to preserve the chronological sequence. This division was critical for evaluating the model's forecasting capabilities using historical data to predict future outcomes. Integrating these pre-processing steps with our data pipeline facilitated compatibility with classification models, thus optimizing the model training and evaluation process.

4.2.3 Classification Model

- Model Selection

  For our binary classification task, we selected Random Forest and XGBoost, known for their efficacy in handling binary outcomes and imbalanced datasets. Random Forest is favored for its interpretability and ability to handle non-linear data through decision trees, reducing overfitting risks. XGBoost excels in efficiency, handling sparse data effectively with its gradient boosting framework, making it ideal for our needs.

- Hyperparameter Tuning

  We utilized Grid Search for systematic tuning of key hyperparameters: n_estimators (number of trees), max_depth (depth of each tree), and learning_rate for XGBoost, alongside n_estimators and max_depth for Random Forest. This method helped us identify the optimal model settings, enhancing predictive accuracy.

  To counter overfitting in XGBoost, we applied L2 (Lambda) and L1 (Alpha) regularization, adjusting penalties on coefficients' magnitude through Grid Search. This step ensured that our model's complexity was balanced with its generalization capability.

- Model Training and Validation

    Employing a time-sensitive dataset split and TimeSeriesSplit for cross-validation, we respected the chronological order of data, essential for accurate future predictions. This setup, along with a 5-fold time-series cross-validation, allowed us to evaluate model performance effectively over various periods. combination that results in the best performance on our validation set, effectively balancing the model's complexity with its ability to generalize to new data.

- Model Evaluation

    Evaluating the performance of XGBoost and Random Forest models, particularly in the context of imbalanced datasets, required careful consideration of the metrics used. In our project, we used Recall and F1-score as a critical metric due to the imbalance in our dataset. Recall measures the proportion of actual positive/negative cases that the model correctly identifies. The formula for recall is below.

    $$Negative\ Recall\ = \frac{True\ Negatives}{True\ Negatives\ +\ False\ Positives}$$

    The F1-score is a harmonic mean of precision and recall, providing a single metric to balance both the precision (the proportion of positive identifications that were actually correct) and the recall of the model.

    $$F1 - Score\ =\ 2\ x\ \frac{Precision\ x\ Recall}{Precision\ +\ Recall}$$

4.2.4 Classification Results

The evaluation of our model using XGBoost and Random Forest yielded the results shown in Figure 7.

Figure 7

Confusion Matrix for the (a) XGBoost Model and (b) Random Forest Model

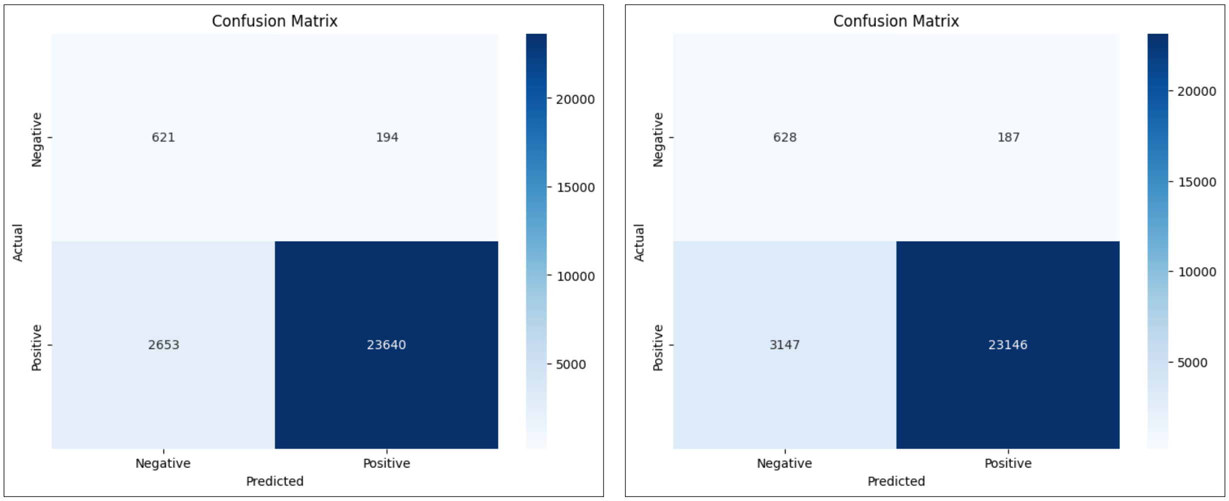

In the confusion matrix, the terms "True Positives," "True Negatives," "False Positives," and "False Negatives" are designated as follows:

- **True Positives (TP):** These are the cases where the model correctly predicts a positive outcome. In the context of the labeling logic, this would be non-stockout cases that were correctly identified as non-stock out. Here, TP is 23,640 for XGBoost and 23,146 for Random Forest.
- **True Negatives (TN):** These are the cases where the model correctly predicts a negative outcome, meaning the stockout situations that were correctly identified as stock outs. Here, TN is 621 for XGBoost and 628 for Random Forest.
- **False Positives (FP):** These occur when the model incorrectly predicts a positive outcome when it is actually negative. It would be a situation where the model predicted non-stockout, but it was actually a stock out. Here, FP is 194 for XGBoost and 187 for Random Forest.
- **False Negatives (FN):** These occur when the model incorrectly predicts a negative outcome when it is actually positive. It would be a situation where the model predicted a stock out, but it was actually non-stock out. Here, FN is 2,653 for XGBoost and 3,147 for Random Forest.

Our classification models demonstrate robust predictive capabilities, with a primary focus on minimizing stockout events. Each model's performance is quantified by the following metrics:

- Positive Recall (Sensitivity): The model accurately predicts non-stockout events a high percentage of the time. This high recall rate indicates strong performance in ensuring inventory levels are deemed sufficient.
- Negative Recall (Specificity): This metric reflects the model's accuracy in predicting actual stockouts.
- Positive F1-Score: The balance between precision and recall for non-stockout predictions is reflected in an F1-score. This confirms the model's precision in predicting when inventory levels are adequate.
- Negative F1-Score: Conversely, the negative F1-score, assessing stockout predictions, highlights the model's performance in this area.

Those metrics, along with their values, are reported in the classification report for both models shown in Table 1.

Table 1

Classification Report Results

Metric	XGBoost		Random Forest	
	Positive	Negative	Positive	Negative
Precision	99%	19%	99%	17%
Recall	90%	76%	88%	77%
F1-score	94%	30%	93%	27%
Overall Accuracy	89%		88%	

In determining the selected model, we chose to prioritize negative recall as the most important metric. This metric reflects the model's ability to accurately predict stockout events. Stockout events are critical and represent minority instances in the dataset; therefore, a higher recall for these events is indicative of better performance by the model in predicting scenarios that are crucial yet less frequent. Given that the Random Forest model demonstrates a negative recall of 77%, compared to the 76% of XGBoost, it is more effective in identifying true stockout events. This higher performance in negative recall is significant for our objectives, leading us to select the Random Forest model for this project.

4.3 Feature Importance Analysis

After training our model, we conducted a feature importance analysis to determine which variables most significantly influence the predictions. This step was vital for understanding the dynamics that affect demand forecasting and inventory management, offering insights that could inform future data collection and feature engineering efforts. For this analysis, we employed SHAP (SHapley Additive exPlanations), a cutting-edge approach that provides a detailed and interpretable explanation of the contribution of each feature to the model's predictions. SHAP values, based on the concept of game theory, attribute the prediction output to each feature's contribution, offering a fair distribution of the prediction among all features.

Using SHAP, we identified not only the most influential features but also the direction and magnitude of their impact on the model's predictions. This granular insight into feature importance helped us understand the complex relationships and interactions between features in our dataset, guiding strategic decisions regarding inventory management. For instance, if SHAP analysis revealed that promotional activities or certain economic indicators significantly influence demand predictions, we adjusted our inventory strategies accordingly to better align with these findings.

Figure 8

SHAP Summary Plot

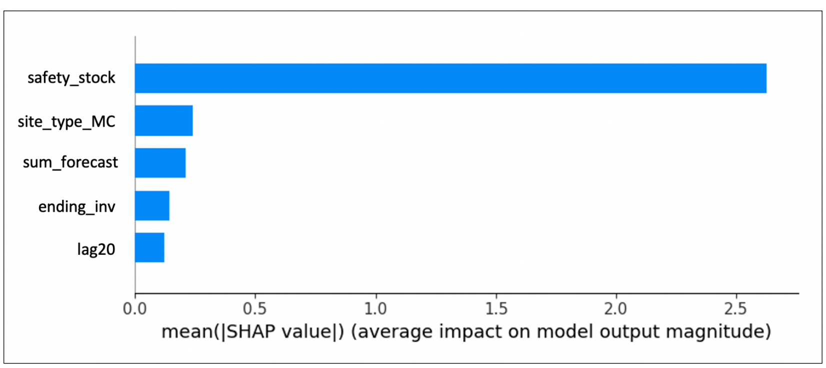

Figure 8 presents the SHAP summary plot, which highlights the five most influential features that impact the model's predictive accuracy for stockout and non-stockout events. This graph clearly indicates that 'safety_stock' is the dominant feature influencing the predictions. In the subsequent step, we explored opportunities to optimize the target inventory levels. This step involved fine-tuning the safety stock across stock-keeping units (SKUs) and sites.

5. Results & Managerial Insights

In this chapter, we present the results of our sensitivity analysis, focusing on its implications for business operations, particularly inventory management. Next, we examine specific SKUs and facility sites identified in the sensitivity analysis, which have the potential for improvement. This examination includes a detailed discussion of the "landslide effect" observed in inventory levels due to seasonal transitions. We conclude with our proposed recommendations for setting target inventory levels. These recommendations were designed to minimize excessive safety stock while simultaneously enhancing service levels, thereby optimizing overall operational efficiency.

5.1 Sensitivity Analysis

The objective of our sensitivity analysis was to evaluate the robustness of our inventory model against demand uncertainty under varying safety stock scenarios. This analysis was crucial for refining inventory strategies by identifying the optimal balance between minimizing stockouts and reducing inventory holding costs.

Our approach involved examining scenarios where the initial model prediction indicated no stockout (positive label where Demand <= Safety Stock + Forecast). Keeping other feature values constant, we systematically adjusted the safety stock levels. These adjustments included reductions in increments of 5%, 10%, 15%, up to 30%, as well as increases by 5%, 10%, up to 30%. We then applied these scenarios to the pre-trained XGBoost model to predict the likelihood of stockouts.

We measured sensitivity by observing changes in the service level, defined as the proportion of non-stockout instances (positive label) relative to the total instances. This metric served as our primary indicator of how sensitive service levels are to changes in safety stock. SKUs and sites where minor reductions in safety stock resulted in only slight declines in service levels were identified as candidates for inventory optimization. These results suggest that safety stock could be reduced without significantly increasing the risk of stockouts. Conversely, SKUs and sites where small increases in safety stock significantly improved service levels were marked as critical for service level optimization.

5.1.1 Results of Sensitivity Analysis

The sensitivity analysis yielded the results shown in Figure 9.

Figure 9

Sensitivity Analysis Result (value: # sku_site)

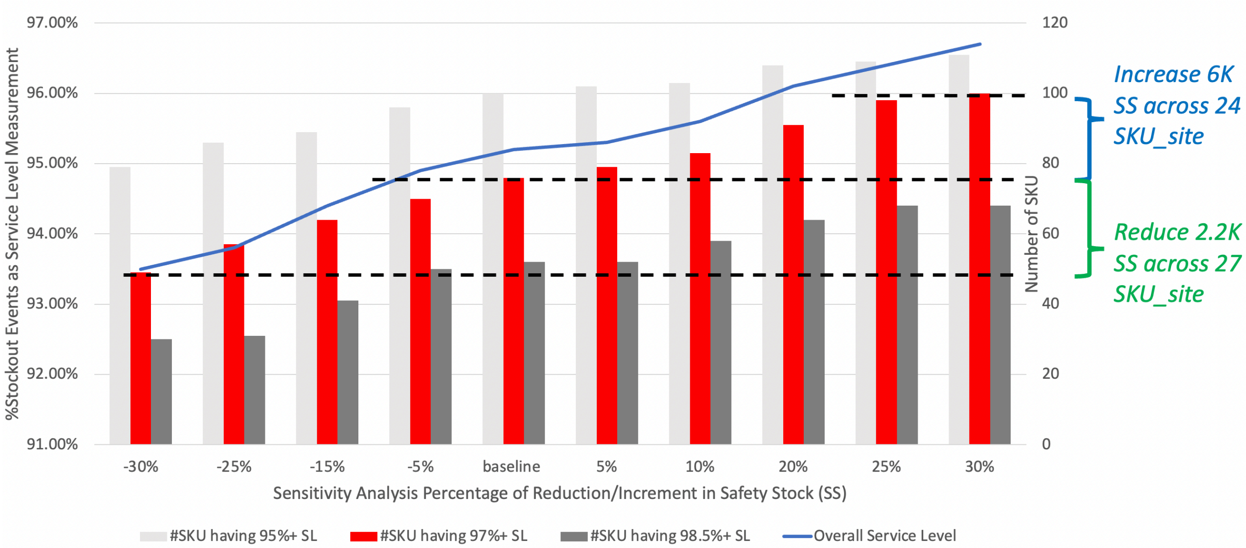

The sensitivity analysis reveals compelling insights into the resilience of the inventory model and its response to varying safety stock levels. At a reduction of 30% in safety stock, the overall service level drops to 93.5%, with significant decreases in the number of SKU_sites maintaining service levels above 95%, 97%, and 98.5%. Conversely, as safety stock increases, there is a noticeable improvement in service levels. At a 30% increase in safety stock, the overall service level peaks at 96.7%, and the number of SKU_sites achieving higher service thresholds also rises, with those maintaining above a 98.5% service level reaching 68, doubling the baseline figure.

These findings underscore the impact of safety stock levels on service performance. It is clear that while reductions in safety stock can lead to cost savings, they also heighten the risk of service level degradation. This is particularly evident as the number of SKU_sites with over 95% service level falls below the baseline when safety stock is decreased. On the flip side, increasing safety stock has a substantial effect on achieving service excellence, with incremental improvements as safety stock levels rise. For instance, there is an increase in SKU_sites with a service level above 97% and 98.5% with each step-up in safety stock, showcasing a positive correlation between safety stock increments and service level enhancements. The analysis presents a strong case for strategic safety stock management as a lever for service level optimization.

5.1.2 Managerial Insights

Our sensitivity analysis facilitated the identification of key inventory adjustments to enhance service levels without disclosing specific SKU_sites, preserving confidentiality. We determined that increasing safety stock by 6,000 units, spread across a select group of 24 locations, could elevate the overall service level to 96.7%. This strategic augmentation of safety stock, executed across carefully chosen points in the supply chain, targets areas where a slight buffer can significantly bolster service robustness.

Conversely, the analysis pinpointed 27 locations where the current safety stock may exceed actual requirements, suggesting an opportunity for inventory right-sizing. By aligning safety stock more closely with demand patterns, the company stands to decrease carrying costs and reduce the likelihood of overstock scenarios, thereby optimizing inventory levels. This balance between increasing safety stock in critical areas and reducing excess in others is key to achieving a lean and responsive supply chain that aligns with market demands and operational goals.

5.2 Revelation of "Landslide Effect"

The analysis of inventory fluctuations throughout seasonal transitions led us to identify a phenomenon we have termed the "landslide effect." This effect poses a challenge for the sponsoring company, as it experiences inventory imbalances during seasonal shifts. Specifically, as the company moves from a low to a high demand season, it finds itself with an excess of inventory. Conversely, when transitioning from a high to a low season, there is a pronounced decrease in service levels due to insufficient inventory (see Figure 10).

Figure 10

Illustration of Landslide Effect

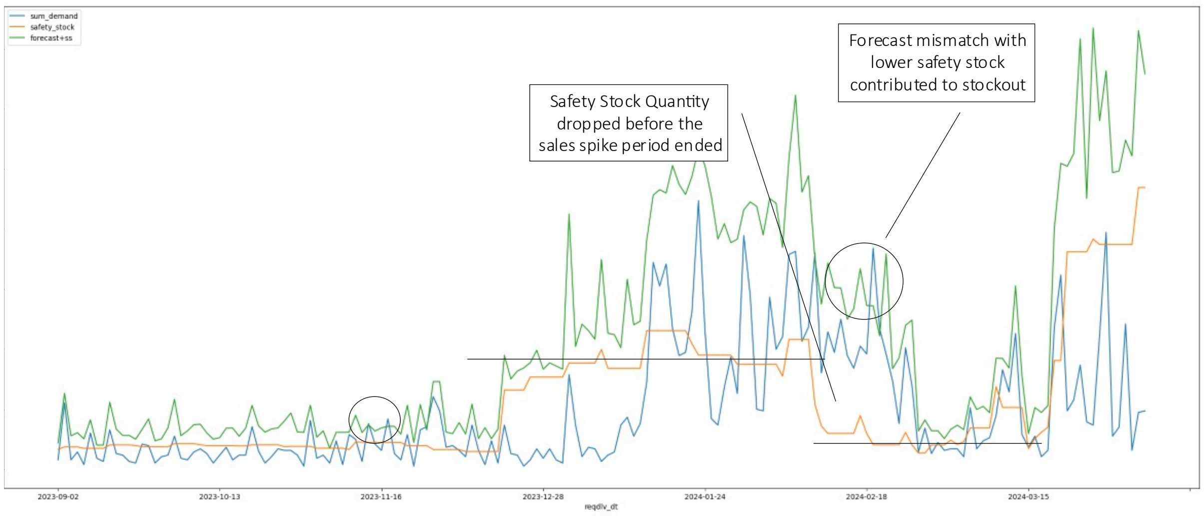

This landslide effect stems from a misalignment in the company's inventory model, which is traditionally backward-looking, basing safety stock targets on past demand over the replenishment lead time. However, the company's current method, which is forward-looking, sets safety stock targets based on anticipated future demand. This discrepancy leads to premature adjustments in inventory levels. During the high season, the forward-coverage, Days of Supply (DOS) model starts to reduce safety stock too early, as its projections begin to include more of the low season's lower demand. This results in inadequate inventory levels while still in the peak season, leading to suboptimal service outcomes.

Furthermore, the forward-coverage DOS approach does not sufficiently consider the trade-offs between inventory savings and the cost implications of stockouts. Without a nuanced evaluation, inventory levels are adjusted without strategic foresight. Adding to the complexity, many practitioners struggle to accurately quantify the cost of stockouts and typically default to setting high service level targets, often in collaboration with supply chain partners. These targets tend to remain fixed across different seasons, which inadvertently contributes to the landslide effect. This outcome, clearly undesirable, necessitates a refined approach to setting inventory targets—one that remains sensitive to seasonal demand fluctuations and the associated service level implications.

5.3 Recommended Approach to Set Target Inventory Level

To address the challenges presented by the landslide effect, we propose an alternative approach for the sponsoring company to set target inventory levels that leverages the appropriate nonstationary calculations for safety stock targets in their planning systems. While a sophisticated solution may involve complex mathematical modeling, a more straightforward, yet effective, strategy is available by applying a Days of Supply (DOS) methodology with some modifications. This practical solution consists of the following steps:

1. **Average Forecast Calculation**: Begin by calculating the average demand forecast for a future period starting from time $t+1$ to $t+T$, where t signifies a given time period and T denotes the replenishment lead time.
2. **DOS Conversion to Safety Stock Units**: Utilize this average forecast to translate the DOS target into a quantifiable safety stock unit target for the period $t+T$.
3. **Planned-Order Release Determination**: Employ the safety stock unit target to ascertain the planned-order release at time t.

Adopting this strategy, although not strictly optimal during seasonal transitions, significantly mitigates the misalignments that arise from a forward-coverage DOS method. While there might be minor discrepancies during these transition phases, they are generally less consequential than those produced by the current approach.

Implementing this new methodology necessitates a shift in perspective for practitioners. It is important to recognize that while the current forward-coverage approach may still be relevant for reporting inventory metrics, it should be adjusted for calculating inventory targets. The recommended approach aligns more closely with the cyclic nature of demand, smoothing the transition between seasons and maintaining service levels without succumbing to the pitfalls of the landslide effect.

5.4 Limitations

In this capstone project, the focus was specifically narrowed down to a particular brand and size within the entire range of products. Across all products, the current business model utilizes the Forward Coverage Inventory Model, assuming that the days of forward coverage are equal to the replenishment or delivery lead time, which is considered a fixed value. It is important to note that this study does not consider lead time variability. Additionally, the multi-echelon inventory optimization problem is not addressed in our research; instead, our analysis concentrates on two layers of distribution. These limitations provide some avenues for future work or the company needs to consider these when implementing the models

6. Conclusion

This capstone project optimizes target inventory level for a leading global Fast Moving Consumer Goods (FMCG) company. Faced with the dual challenge of preventing stockouts and minimizing the costs associated with excess inventory, this project employed machine learning techniques to set optimal target inventory levels, enhancing overall service levels.

Our primary deliverable was a robust machine learning model that accurately predicts stockout events for each stock keeping unit (SKU) and warehouse (site) combination. Achieving an impressive 91% accuracy, evaluated by f1-score and recall, this model leverages Random Forest technique. Furthermore, we conducted comprehensive sensitivity analyses to understand the correlations between service levels and varying inventory scenarios, leading to substantial strategic insights.

The sensitivity analysis unveiled significant opportunities for inventory optimization by identifying potential overstocks and understocks across various SKU_sites. We discovered the "landslide effect," a critical challenge during seasonal transitions that leads to mismatches in inventory levels. This phenomenon was effectively addressed by a new inventory setting approach we proposed, which focuses on backward demand data instead of future forecast to set safety stock targets, steering clear of the intrinsic flaw of forward-looking inventory method.

Moving forward, we recommend that the sponsoring company adopt the modified Days of Supply (DOS) approach to recalibrate safety stock levels based on retrospective demand forecasts rather than prospective estimates. This shift will help in smoothing out the inventory levels throughout the seasonal cycles, thereby maintaining higher service levels and reducing unnecessary stock holding costs. Continuous improvement of the machine learning models and further exploration into predictive analytics are suggested to keep enhancing the accuracy of stockout predictions and overall inventory management efficiency.

This project not only achieved its aim of refining inventory management practices but also provided a scalable model for the FMCG industry to follow, combining advanced analytical techniques with practical inventory management solutions to tackle complex supply chain challenges.

REFERENCES

Ban, G.-Y., & Rudin, C. (2019). The Big Data Newsvendor: Practical Insights from Machine Learning. *Operations Research*, *67*(1), 90–108. https://doi.org/10.1287/opre.2018.1757

Bollapragada, S., & Morton, T. E. (1999). A Simple Heuristic for Computing Nonstationary (s, S) Policies. *Operations Research*, *47*(4), 576–584. https://doi.org/10.1287/opre.47.4.576

Bureau of Labor Statistics. (2023a). *Consumer Price Index Archived News Releases: U.S. Bureau of Labor Statistics*. https://www.bls.gov/bls/news-release/cpi.htm#2018

Bureau of Labor Statistics. (2023b). *CPI Handbook of Methods Appendix 2: Content of CPI Entry Level Items : U.S. Bureau of Labor Statistics*. https://www.bls.gov/cpi/additional-resources/entry-level-item-descriptions.htm

Bureau of Labor Statistics. (2023c). *Overview: Handbook of Methods: U.S. Bureau of Labor Statistics*. https://www.bls.gov/opub/hom/cpi/home.htm

Carbonneau, R., Laframboise, K., & Vahidov, R. (2008). Application of machine learning techniques for supply chain demand forecasting. *European Journal of Operational Research*, *184*(3), 1140–1154. https://doi.org/10.1016/j.ejor.2006.12.004

Carvalho, H., Naghshineh, B., Govindan, K., & Cruz-Machado, V. (2022). The resilience of on-time delivery to capacity and material shortages: An empirical investigation in the automotive supply chain. *Computers & Industrial Engineering*, *171*, 108375. https://doi.org/10.1016/j.cie.2022.108375

Chopra, S., & Meindl, P. (2016). *Supply chain management: Strategy, planning, and operation* (Sixth Edition). Pearson.

Clausen, J. B. B., & Li, H. (2022). Big data driven order-up-to level model: Application of machine learning. *Computers & Operations Research*, *139*, 105641. https://doi.org/10.1016/j.cor.2021.105641

De Melo, Á. S. (2019). A machine learning approach to the optimization of inventory management policies.

Gallegoi, G., & Moon, I. (2023). *The Distribution Free Newsboy Problem: Review and Extensions*. *44*(8).

Gardner, E. S. (1990). Evaluating Forecast Performance in an Inventory Control System. *Management Science*, *36*(4), 490–499. https://doi.org/10.1287/mnsc.36.4.490

Giannoccaro, I., & Pontrandolfo, P. (2002). Inventory management in supply chains: A reinforcement learning approach. *International Journal of Production Economics*, *78*(2), 153–161. https://doi.org/10.1016/S0925-5273(00)00156-0

Gumus, M., & Kiran, M. S. (2017). Crude oil price forecasting using XGBoost. *2017 International Conference on Computer Science and Engineering (UBMK)*, 1100–1103. https://doi.org/10.1109/UBMK.2017.8093500

Gundogdu, B., & Maloney, J. (2019). *Comparison and Financial Assessment of Demand Forecasting Methodologies for Seasonal CPGs*.

Guo, Z. X., Wong, W. K., & Li, M. (2013). A multivariate intelligent decision-making model for retail sales forecasting. *Decision Support Systems*, *55*(1), 247–255. https://doi.org/10.1016/j.dss.2013.01.026

Hartunian, G. (2015, April 12). *Four Steps to Service Level-Driven Inventory Planning*. https://www.scmr.com/article/four_steps_to_service_level_driven_inventory_planning

Huber, J., Müller, S., Fleischmann, M., & Stuckenschmidt, H. (2019). A data-driven newsvendor problem: From data to decision. *European Journal of Operational Research*, *278*(3), 904–915. https://doi.org/10.1016/j.ejor.2019.04.043

Huh, W. T., & Rusmevichientong, P. (2009). A Nonparametric Asymptotic Analysis of Inventory Planning with Censored Demand. *Mathematics of Operations Research*, *34*(1), 103–123. https://doi.org/10.1287/moor.1080.0355

Hyndman, R. J., & Koehler, A. B. (2006). Another look at measures of forecast accuracy. *International Journal of Forecasting*, *22*(4), 679–688. https://doi.org/10.1016/j.ijforecast.2006.03.001

Inderfurth, K. (1991). Safety stock optimization in multi-stage inventory systems. *International Journal of Production Economics*, *24*(1–2), 103–113. https://doi.org/10.1016/0925-5273(91)90157-O

ITC Infotech. (2020). *Inventory Management and Optimization for an FMCG Manufacturing Company*. https://www.anylogic.com/resources/case-studies/inventory-management-and-optimization-for-an-fmcg-manufacturing-company/

Kara, A., & Dogan, I. (2018). Reinforcement learning approaches for specifying ordering policies of perishable inventory systems. *Expert Systems with Applications*, *91*, 150–158. https://doi.org/10.1016/j.eswa.2017.08.046

Khademi Zare, H., Yahia Mehrjerdi, Y., Akhbari Varkani, M., & Makui, A. (2014). *A novel continuous knn prediction algorithm to improve manufacturing policies in a vmi supply chain*. International Journal of Engineering, 27(11), 1681-1690. https://doi.org/10.5829/idosi.ije.2014.27.11b.05

Kotler, P., & Armstrong, G. (2014). *Principles of Marketing (15th ed.)*. Pearson.

Lee, M., & Siddiqui, K. (2023). *Case Fill Rate Prediction*. Department of Supply Chain Management, Massachusetts Institute of Technology.

Lohman, C., Fortuin, L., & Wouters, M. (2004). Designing a performance measurement system: A case study. *European Journal of Operational Research*, *156*(2), 267–286. https://doi.org/10.1016/S0377-2217(02)00918-9

Maheshwari, S., Gautam, P., & Jaggi, C. K. (2021). *Role of Big Data Analytics in supply chain management: current trends and future perspectives.* International Journal of Production Research, 59(6), 1875-1900.

Michalski et al. - 1983—*Machine Learning.pdf*. (n.d.).

Namir, K., Labriji, H., & Ben Lahmar, E. H. (2022). Decision Support Tool for Dynamic Inventory Management using Machine Learning, Time Series and Combinatorial Optimization. *Procedia Computer Science*, *198*, 423–428. https://doi.org/10.1016/j.procs.2021.12.264

Perakis and Roels—2008—Regret in the Newsvendor Model with Partial Inform.pdf. (n.d.).

Perez, H. D., Hubbs, C. D., Li, C., & Grossmann, I. E. (2021). Algorithmic Approaches to Inventory Management Optimization. *Processes*, *9*(1), 102. https://doi.org/10.3390/pr9010102

Priore, P., Ponte, B., Rosillo, R., & De La Fuente, D. (2019). Applying machine learning to the dynamic selection of replenishment policies in fast-changing supply chain environments. *International Journal of Production Research*, *57*(11), 3663–3677. https://doi.org/10.1080/00207543.2018.1552369

Rapidor. (2023, August 30). *Inventory Management Techniques for FMCG*. https://www.rapidor.co/inventory-management-techniques-for-fmcg/

Seyedan, M., Mafakheri, F., & Wang, C. (2023). Order-up-to-level inventory optimization model using time-series demand forecasting with ensemble deep learning. *Supply Chain Analytics*, *3*, 100024. https://doi.org/10.1016/j.sca.2023.100024

Silver, E. A., Pyke, D. F., Silver, E. A., & Silver, E. A. (2017). *Inventory and production management in supply chains* (Fourth Edition). CRC Press, Taylor & Francis Group.

Teunter, R. H., Syntetos, A. A., & Babai, M. Z. (2017). Stock keeping unit fill rate specification. *European Journal of Operational Research*, *259*(3), 917–925. https://doi.org/10.1016/j.ejor.2016.11.017

Vidal, G. H. (2023). Deterministic and Stochastic İnventory Models in Production Systems: A Review of the Literature. *Process Integration and Optimization for Sustainability*, *7*(1–2), 29–50. https://doi.org/10.1007/s41660-022-00299-3

Wang, J., Li, X., & Zhu, X. (2012). Intelligent dynamic control of stochastic economic lot scheduling by agent-based reinforcement learning. *International Journal of Production Research*, *50*(16), 4381–4395. https://doi.org/10.1080/00207543.2011.592158

Zhang, Y., & Yang, X. (2016). Online ordering policies for a two-product, multi-period stationary newsvendor problem. *Computers & Operations Research*, *74*, 143–151. https://doi.org/10.1016/j.cor.2016.04.031