Smart Stock: Demand-Driven Inventory for Lower Costs and Higher Returns

by

Natalia Higdon
MBA, University of Houston-Victoria, 2016

and

Sreerag Gangadharan
MBA, Liverpool John Moores University, 2010

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

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

May 2025

Signature of Author: ____

Natalia Higdon
Department of Supply Chain Management
May 9, 2025

Signature of Author: ____

Sreerag Gangadharan
Department of Supply Chain Management
May 9, 2025

Certified by: ____

Ilya Jackson
Research Scientist
Capstone Advisor

Certified by: ____

Jaime E. Macias Aguayo
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

Smart Stock: Demand-Driven Inventory for Lower Costs and Higher Returns

by

Natalia Higdon

and

Sreerag Gangadharan

Submitted to the Program in Supply Chain Management
on May 9, 2025 in Partial Fulfillment of the
Requirements for the Degree of Master of Applied Science in Supply Chain Management

ABSTRACT

The energy sector faces persistent challenges in managing spare parts inventory due to demand volatility, operational complexity, and capital intensity. Our Sponsor, one of the world's largest energy companies, sought to optimize inventory by aligning stock levels more closely with business demand, with a goal of reducing excess holdings and improving service reliability. To address this challenge, we conducted a comprehensive analysis of over 40,000 materials across nine business units, segmenting them into four categories based on demand intermittence and variability metrics. We then tested seven machine learning models across the categories identified, ultimately selecting the two best-performing ones. For materials lacking sufficient historical data required for machine learning, we implemented a pragmatic alternative using average consumption. This two-pronged approach revealed a nearly 27% inventory reduction opportunity while achieving a 95% target service level. Our findings demonstrate that a dynamic, demand-driven inventory policy can unlock substantial working capital, improve asset readiness, and enhance operational resilience, providing a scalable framework for enterprise-wide inventory optimization.

Capstone Advisor: Ilya Jackson
Title: Research Scientist

Capstone Co-Advisor: Jaime E Macias Aguayo
Title: Postdoctoral Associate

ACKNOWLEDGMENTS

We would like to express our deepest gratitude to our Capstone advisors, Ilya Jackson and Jaime Macias Aguayo, for their invaluable guidance, thoughtful feedback, and steady support throughout this project. Their expertise and encouragement challenged us to refine our ideas and continuously elevate the quality of our work.

We are immensely grateful to the sponsoring company for the opportunity to work on such an impactful and relevant problem. In particular, we thank the Procurement and Materials Management leader for the generous feedback, strategic insights, and ongoing engagement throughout the project.

We also extend sincere thanks to Toby Gooley, our writing coach, whose mentorship helped us refine our narrative with clarity, structure, and purpose.

Finally, we thank the MIT SCM community—our classmates, faculty, and staff—for fostering an environment of collaboration, intellectual curiosity, and shared inspiration.

TABLE OF CONTENTS

1. INTRODUCTION

Advancements in energy production have been a catalyst for global economic growth and scientific discovery over the past century. Access to reliable and affordable energy has elevated living standards, strengthened public health and education, and expanded economic opportunity. Today, the evolution of energy infrastructure stands not only as a foundation of societal advancement but as a driving force for innovation worldwide.

According to the Energy Institute's latest Statistical Review of World Energy (2024), over 80% of the world's energy still comes from fossil fuels. However, renewable energy sources such as wind, solar, and hydroelectric power are expanding rapidly as the world works toward a more sustainable future. Amid this transition, one fact is clear: energy production, whether from oil and gas, wind, solar, or nuclear, remains highly capital-intensive, with global energy investments exceeding USD 3 trillion for the first time in 2024 (International Energy Agency [IEA], 2024).

While energy infrastructure development requires substantial investments, maintaining the operational integrity of production systems is equally critical and costly. In this project, we focus specifically on how advancements in machine learning for materials management applications can support energy production, where asset uptime and equipment reliability are key.

1.1 MOTIVATION

Demand forecasting for spare parts has historically posed a major issue for the oil & gas industry, driven by market volatility, geopolitical forces, and a lack of standardization. Hydrocarbon extraction, transportation, and refining methods vary depending on a range of variables, including the type of oil, the gases present in the reservoir, geographic area, onshore or offshore environment, pipeline availability and ownership, and the refining footprint. These operational variations are necessary due to the considerable structural and geographical diversity of hydrocarbons (Wilkes & Schwarzbauer, 2010). The combination of extreme operating conditions, complex processes, and stringent regulatory requirements makes oil & gas one of the most capital-intensive sectors of the world economy, requiring a wide range of generic and specialized equipment, which is hard to plan for and expensive to procure, maintain, and store. Furthermore, given that oil & gas companies operate in commodity markets characterized by high fixed costs, minimal pricing power, and continuous investment needs, the efficiency of production processes, manufacturing capabilities, and optimized workflows becomes a strategic source of competitive advantage (Moser, Isaksson, et al., 2021). The need to enhance supply chain workflow automation required for servicing the industry in such an environment served as a central motivation for this Capstone project.

The COVID-19 pandemic led to unprecedented declines in oil demand due to widespread lockdowns and travel restrictions. In early April 2020, global oil demand dropped by over 30 million barrels per day, representing a 33% contraction in demand (Le et al., 2021). This sharp decline in consumption contributed to a historic plunge in oil prices, with West Texas Intermediate (WTI) crude oil futures falling below zero for the first time in history on April 20, 2020. Compounding these challenges, the availability of materials and services essential for maintaining operating assets was severely disrupted. The industry's supply chains faced significant interruptions, including shortages of critical components, delays in equipment deliveries, and workforce constraints, all of which negatively impacted routine maintenance and operational continuity (Piya et al., 2022).

Our Sponsor, one of the world's largest oil & gas companies, exemplifies the challenges posed by industry dynamics, pandemic disruptions, and forecasting uncertainty in this complex operational environment. Headquartered in Texas, U.S.A., our Sponsor is a multinational, publicly traded company with an annual production of millions of barrels of oil equivalent per day (Sponsor company, 2024).

Prompted by the COVID-19 pandemic, our Sponsor organization conducted a series of internal evaluations and consultant-led assessments focused on inventory management opportunities. These assessments revealed that the Company had been challenged with inventory stockout events in some cases and inventory surplus in others. The root cause of these challenges was attributed to a misalignment of demand and supply flows due to the lack of integration between operations demand planning, and spare parts inventory. For the purposes of this study, the terms *spare parts*, *safety stock*, *inventory*, and *critical spares* will be used interchangeably, consistent with both industry standards and the terminology reviewed in the State of Practice (Section 2).

The misalignment of business demand and available spare parts inventory has led to elevated costs and supply assurance issues, which have negatively impacted financial performance in both the upstream (exploration and production) and downstream (refining) sectors, leading to underperformance relative to industry peers (Sponsor Company, personal communication, 2024).

1.2 PROBLEM STATEMENT

Demand for spare parts is usually prompted by routine maintenance work, equipment failure, or component failure, which means that demand patterns are volatile and can be unpredictable. To ensure the availability of spare parts needed to maintain operations and perform the necessary repairs, most oil & gas companies, including our Sponsor, maintain sizeable inventory levels, which go beyond supporting the immediate business needs.

Our Sponsor Organization procures and maintains spare parts to service four upstream and five downstream Business Units. These parts are stored in 17 major warehouse hubs and 30 smaller warehouse spokes across the U.S., which is a complex network of storage locations.

The Sponsor's key objective is to reduce inventory holding costs while meeting business requirements, defined through Service Level Agreements (SLAs). To support this goal, our project focuses on aggregating demand signals, identifying underlying usage patterns, and incorporating business priorities to establish demand-driven inventory policies tailored to the facilities in scope. This approach aims to provide actionable insights to help balance inventory availability, service levels, and financial investment.

Accordingly, we propose the following questions:

1. How can the Sponsor improve inventory availability while reducing inventory costs?
2. What are the opportunities to improve the current approach to demand forecasting and inventory management?
3. What is the overall value of introducing a dynamic, demand-driven inventory policy?

1.3 SCOPE: PROJECT GOALS AND EXPECTED OUTCOMES

The project's objective is to provide the Sponsor with a recommendation for a demand-driven inventory policy that will reduce stock levels, lower inventory holding costs, and improve service levels. In addition to quantitative optimization, the project team will provide actionable qualitative interpretation by leveraging the team's extensive oil & gas SCM experience.

The deliverables to the Sponsor will include:

1. A methodology for defining a data-driven spare parts demand plan.
2. A methodology for defining a dynamic, demand-driven inventory policy.

3. An in-depth assessment of the current inventory state, with tailored recommendations for rebalancing stock levels to meet a 95% target service level, as defined by the Sponsor.

2. STATE OF THE PRACTICE

Oil & gas is a 24/7 operation. In the field, three things matter most: people's safety, full containment of hydrocarbons, and their constant flow. Ensuring the availability of the right spare parts at the right time and place is essential to achieving these goals and avoiding costly downtime. To reduce the risk of stockouts, operators often over-purchase spare parts, leading to excess inventory and elevated holding costs—an issue our Sponsor organization also faces. Addressing this challenge requires a shift toward more precise, demand-driven inventory management, informed by improved forecasting and adjusted for material criticality and service level targets. In this section, we examine industry best practices and academic insights in three key areas: demand forecasting and safety stock methodologies, common challenges in inventory management for asset-intensive operations, and product segmentation techniques that support spare parts planning and replenishment.

2.1 DEMAND FORECASTING AND SAFETY STOCK

Demand forecasting is a well-established discipline at both the strategic and operational levels. Syntetos et al. (2016) highlight significant advances in this area, encompassing statistical, judgmental, and hybrid forecasting techniques. At its core, demand forecasting aims to minimize costs and maximize service by predicting future requirements based on historical usage, market signals, and internal and external variables (Tadayonrad & Ndiaye, 2023). Demand forecasting goes hand in hand with safety stock determination, which is used to soften the impact of unplanned changes in supply and demand. Establishing appropriate safety stock levels is critical for companies managing physical assets, as it enables them to meet service level targets while minimizing inventory holding costs. This plays a crucial role in meeting consumer demand and production goals and helping to maintain the integrity of companies' Property, Plant & Equipment (PP&E).

In the oil & gas industry, determining optimal stock levels for PP&E maintenance and operations is particularly challenging. A substantial portion of spare parts demand arises from equipment and component failures, which are inherently unpredictable. Furthermore, the diversity of production facilities—stemming from a history of acquisitions and decentralized decision-making—results in varied equipment configurations across sites. This variability contributes to a wide spectrum of demand patterns, ranging from frequent and predictable to highly volatile and infrequent.

Understanding the demand characteristics of different material types is essential for designing an effective inventory policy that balances inventory availability with reduced management costs. To aid with this, we will introduce a series of segmentation methods to classify and cluster materials with similar characteristics. These methods will form the foundation for a targeted approach to inventory management and will be discussed in detail in Section 2.3: Product Segmentation for Demand Forecasting.

2.2 INVENTORY MANAGEMENT OVERVIEW AND COMMON CHALLENGES

As a discipline, inventory management focuses on optimizing inventory levels while reducing costs and meeting established service levels, which is a challenge. Basten & Houtum (2014) point out that maintaining optimal inventory levels is of high importance, and cost minimization must be balanced with capital expenditures needed to acquire and maintain adequate stock levels. To optimize inventory investments and ensure customer satisfaction, a company must gain a clear understanding of the linkage between demand forecasting and safety stock levels (Tadayonrad & Ndiaye, 2023).

Another common challenge related to inventory management stems from minimizing surplus to reduce inventory holding costs. In the oil & gas industry, inconsistent and volatile demand inevitably leads to excess inventory, also known as dead stock or surplus. For the purposes of this capstone, we will use these terms interchangeably. Surplus inventory is costly to manage as it takes up valuable shelf space, requiring preservation, active management, and documentation to meet regulatory compliance requirements. Early studies established foundational benchmarks for these costs. Richardson (1995) quantified annual inventory holding costs as a range between 25% and 55% of the material cost. Thummalapalli (2010) refined these estimates to a 20-40% range, reflecting incremental efficiency gains stemming from technological advancements. More recent studies show the range of inventory holding costs remains close to 25% even with the technological advancements (Odedairo, Alaba, & Edem, 2020). But the cost overruns don't stop here. In addition to carrying costs, other burdens are associated with excess stock. From the accounting standpoint, surplus inventory is treated as an asset on a company's balance sheet, which means that it inflates asset value, reducing Return on Investment (ROI) and misrepresenting inventory turnover (Grant, Karagianni, & Li, 2006). We will use these concepts further in our research to quantify the value of inventory optimization.

2.3 PRODUCT SEGMENTATION FOR DEMAND FORECASTING

Integrating demand forecasting with inventory management policies for various inventory types can help effectively manage company stock levels (Bacchetti & Saccani, 2012). Optimal stock levels depend on specific business requirements and underlying demand trends. To recommend an effective inventory policy, it is essential first to analyze and differentiate the demand patterns of various product types. This analysis is often facilitated by segmentation techniques that cluster materials based on their criticality and usage patterns. Several of these techniques have been described and adopted in this Capstone project to inform inventory policy recommendations.

Research into forecasting methods for irregular demand began in the early 1970s, with Croston (1972) introducing a technique specifically designed for intermittent demand. His method separated demand size from the frequency of demand occurrences, offering a practical solution for forecasting low-frequency items. However, the formal categorization of demand patterns as a distinct analytical practice emerged later. Williams (1984) was among the first to propose a structured classification system based on demand regularity, lead time variability, and unpredictability, categorizing products as smooth, sporadic, or slow-moving. Building on this, Gilbert and Finch (1985) introduced the widely adopted ABC classification method, segmenting inventory into A, B, and C categories based on value, criticality, or usage frequency. While still common, ABC analysis lacks the precision needed to account for demand volatility in spare parts forecasting.

Croston's original method was later refined to address key limitations. Rao (1973) introduced a bias correction to reduce the tendency of Croston's method to overestimate demand for highly intermittent items. Syntetos and Boylan (2005) expanded this line of work by proposing a systematic classification framework using two key metrics: intermittence (p), which reflects the average time between non-zero demand occurrences, and coefficient of variation squared (CV^2), which captures demand variability in magnitude. Based on these dimensions, demand can be segmented into four primary categories:

1. Erratic– Characterized by low intermittence (frequent demand occurrences) but high variability in demand magnitude
2. Lumpy– Defined by high intermittence (sporadic demand occurrences) and high variability, leading to unpredictable demand patterns
3. Smooth– Exhibits low intermittence (consistent demand occurrences) and low variability, resulting in a stable demand profile
4. Intermittent– Displays high intermittence (demand occurs irregularly) but low variability, meaning demand sizes remain relatively stable when they do occur.

This classification provides a structured approach to analyzing demand patterns, particularly in industries where forecasting intermittent or highly variable demand is a critical challenge.

Figure 1: Cutoff values (Syntetos-Boylan Approximation of Croston's method)

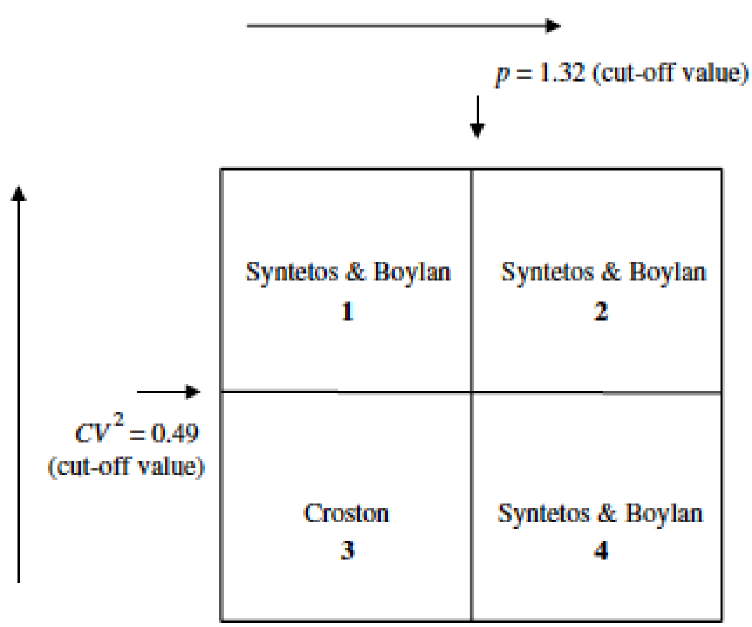

Note. From Syntetos, Boylan & Croston (2005)

The characterization of demand patterns facilitates the selection of forecasting methods, which is an essential element in many of the inventory control software packages (Syntetos, Boylan & Croston, 2005). It is important to note that many companies may use manual inventory estimates or manually adjust software recommendations. After surveying 240 US corporations, Sanders and Manrodt (2003) found that around 60% of the corporations that use forecasting software adjust their forecasts based on human judgment. Citing this study along with others in their research, Perera et al. (2019) argue that human judgment and forecasting are fundamentally inseparable. Even with all the advancements in quantitative forecasting methods, a judgmentally adjusted forecast helps add the qualitative parameters that could not be added in the statistical models. Fildes et al. (2009) noted down their main findings after further research into this area: (a) managerial adjustments do improve accuracy; (b) small adjustments, less than 10%, are not worth making and should be discouraged; and (c) negative adjustments are more effective than positive ones.

Most inventory management software allows companies to consider demand patterns to determine optimal stock levels. Syntetos & Boylan (2005) point out that, in many cases, such characterization is arbitrary and may be further improved by scientifically selecting the forecasting method most appropriate for each type of forecast. In this project, we plot demand intermittence (p) and demand variability (CV^2) to segment inventory into erratic, lumpy, smooth, and intermittent categories. We then adjusted inventory level recommendations for each category to better match material demand patterns.

3 METHODOLOGY

Following a comprehensive literature review and extensive discussions with the Sponsor Company and academic advisors, we focused on the primary objective of this project: developing a scalable, demand-driven inventory policy that optimizes spare parts inventory levels, reduces costs, and improves service reliability. Our methodology was structured to bridge the gap between theoretical best practices in inventory management and the operational realities of a global oil & gas enterprise. Specifically, this chapter outlines the step-by-step approach we developed, from initial scoping through demand segmentation, machine learning, and dynamic inventory policy formulation, leading to actionable recommendations for value capture.

Building on the frameworks reviewed in the State of the Practice (Chapter 2), we adopted several industry-proven methods tailored to the Sponsor's operational environment. We applied demand segmentation techniques based on intermittence and variability metrics (Syntetos, Boylan, & Croston, 2005) to differentiate material demand patterns, enabling targeted forecasting and stocking strategies. For demand forecasting, we evaluated a range of models, including AutoARIMA, Croston's method variations, and newer machine learning approaches, ultimately selecting models best suited for the Sponsor's unique demand profiles. In cases where sufficient historical data was unavailable, we modified the advanced safety stock determination methodologies by incorporating a three-month average consumption to maintain comprehensive inventory coverage. The methodological adaptations described here reflect both industry best practices and practical adjustments needed to address data limitations, operational variability, and system constraints identified through collaboration with the Sponsor organization.

Our approach is structured as follows:

1. Scoping
2. Data Preparation
3. Initial Analysis and Visualization
4. Material Segmentation
5. Machine Learning
6. Demand Forecasting
7. Inventory Level Recommendation
8. Value Interpretation

3.1 SCOPE DEFINITION

Given the scale and complexity of the Sponsor's global operations, clearly defining the project scope was essential to ensure feasibility and focus. Our Capstone focused on the Sponsor's North American upstream and downstream operations, where inventory management for maintenance and operating spares is both critical and cost-intensive. The scope was limited to spare parts stored in company-owned warehouse facilities and managed through the organization's Enterprise Resource Planning (ERP) system. These materials support ongoing maintenance and asset reliability efforts and are distinct from capital project or consumable items, which were excluded from analysis.

3.2 DATA PREPARATION

Data preparation is a critical step in any inventory optimization project. The structure of the data model, data quality, and data completeness directly influence the validity and reliability of the analysis. This section outlines the datasets provided by the Sponsor Company and details the processes applied to transform the raw data into a format suitable for further analysis.

The Sponsor Company supplied two key datasets for this project: the On-Hand Inventory Report and the Historical Inventory Movement Report.

The On-Hand Inventory Report captures a comprehensive snapshot of all inventory stored across the company's North American locations. This dataset includes the following attributes: Business Unit Name, Storage Location (SLOC), Material Number, Material Description, Quantity, Moving Average Price (MAP), and Inventory Value. Each inventory item is unique to the SLOC where it is housed. The Inventory Report provides a complete inventory list for each location rather than a sample, ensuring comprehensive coverage of the Sponsor's inventory. The snapshot of the current inventory level was taken in November 2024.

The Historical Inventory Movement Report provides a detailed record of every inventory movement transaction across the company's North American locations over the past eight years. This dataset includes the following attributes: Business Unit Name, Storage Location (SLOC), Material Number, Material Description, Movement Type, and Quantity per transaction. This dataset offers a comprehensive transaction history spanning from 2017 to 2024, capturing every inventory movement transaction during the specified period. This report provides the information needed to understand the demand patterns, variability, and inventory trends.

A structured series of Python-based data processing steps was applied to transform the raw dataset into a format suitable for analysis, as outlined in Section 3.4: Material Segmentation Methodology. These preprocessing steps involved aggregating transactional records into daily and monthly demand signals, thereby enabling the detection of underlying usage patterns. This transformation was critical to ensure the dataset was appropriately structured to support subsequent segmentation and inventory optimization efforts.

3.3 INITIAL ANALYSIS AND VISUALIZATION

In this step, we explored the data structure, analyzed business units in scope, demand distribution, early trends, detected outliers, and evaluated the shape and scale of the data. Building on the preprocessing steps outlined in Section 3.1 (Data Preparation), this analysis facilitated the identification of demand patterns and the calculation of descriptive statistics, providing critical insights into the dataset.

We began by examining the number of transactions per year across all business units. These transactions include all goods issued out (demand signals), goods returned, warehouse transfers, and scrap movements. Figure 2 illustrates the annual distribution of transactions withing the dataset.

Figure 2: Annual Number of Transactions

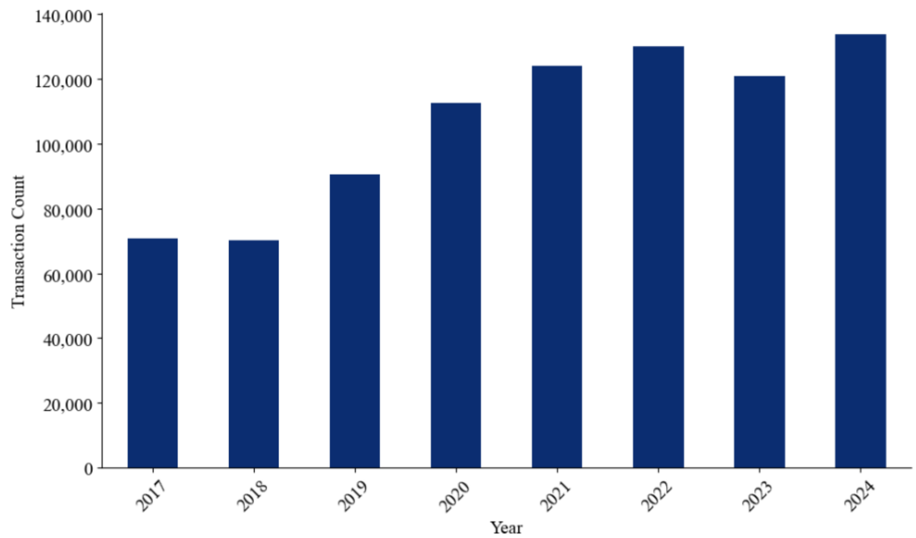

After discussions with the sponsor company, a decision was made to use data from 2017 onward for this Capstone project as outlined in Section 3.2. That data was used for further analysis, as well as for training and testing the machine learning models.

We further filtered the dataset to only include material movements associated with the plans in scope. We consolidated individual transactions at the monthly level to create a demand time series for each material. A monthly cadence was selected due to the sporadic nature of demand for operating spares. After filtering, the dataset included 40,687 unique materials across nine Business Units (BUs). The unique materials count per BU is displayed in Figure 3.

Figure 3: Unique Material Counts per Business Unit

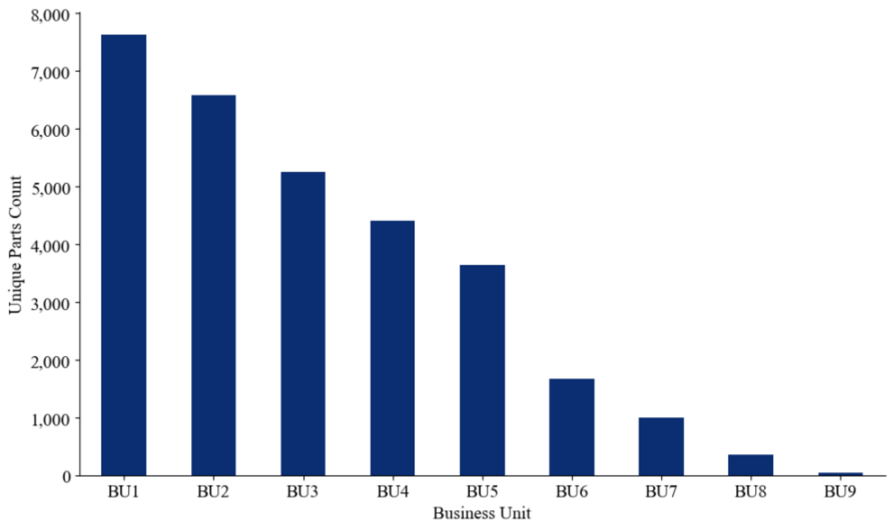

After the data cleansing work, we evaluated the annual transaction counts for the defined scope to identify materials with appropriate characteristics for machine learning. Specifically, we identified materials with adequate training data from 01/01/2022 to 12/31/2023 and testing data from 01/01/2024 to 11/01/2024. Figure 4 below indicates annual transaction counts for materials in scope. These transactions will inform inventory policy recommendations.

Figure 4: Annual Transaction Counts for Materials in Scope

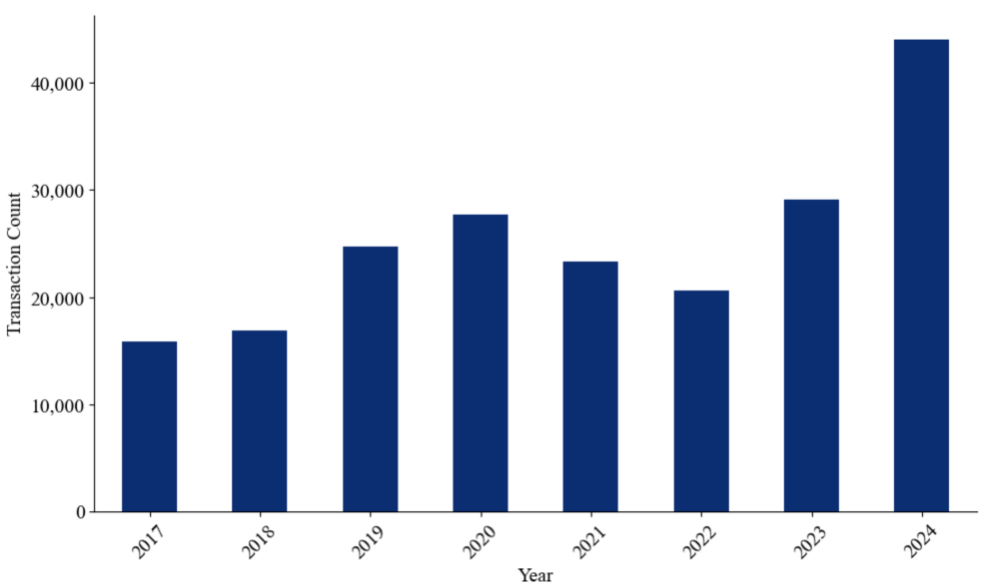

3.4 MATERIAL SEGMENTATION

Our analysis and stakeholder conversations confirmed that demand for materials is not uniform, given that maintenance and operations activities are not always predictable, and material utilization for parts varies depending on application and operating conditions. Given the non-uniform demand for spare parts, not all materials had enough data to perform machine learning for demand forecasting. Of the 40,687 parts, only 1,596 had adequate data coverage for training and testing machine learning models. These parts represented high-volume, high-spend, and high-criticality items that required robust planning and effective inventory management. However, recognizing the value of the broader dataset, analytical methods were applied to calculate the average monthly consumption for materials with insufficient data, ensuring a comprehensive inventory review that could inform the Sponsor's inventory policy.

For spare parts eligible for machine learning, we hypothesized that different demand patterns would require distinct forecasting models, depending on demand variability and the frequency of demand occurrences. To test this hypothesis, we conducted a detailed material segmentation analysis and applied multiple machine learning models to each segment to identify the best-performing forecasting approach.

To categorize SKUs, we analyzed monthly demand patterns and classified them into four categories: smooth, intermittent, erratic, and lumpy. This classification follows the methodology proposed by Syntetos, Boylan, and Croston (2005), as outlined in Chapter 2: State of the Practice. To implement this framework, we computed two key metrics: the squared coefficient of variation (CV^2), which measures demand variability, and the average inter-demand interval (p), which quantifies the time between non-zero demand occurrences. These metrics served as the basis for demand segmentation and informed the selection of forecasting models tailored to each demand profile. Consistent with best practices in the literature, we applied commonly used threshold values of $p = 1.32$ and $CV^2 = 1.0$ to delineate the segments (e.g., Boylan, Syntetos, & Karakostas, 2008; do Rego & de Mesquita, 2015). A higher p indicates lower demand frequency, while a higher CV^2 reflects greater variability in demand magnitude (Chuang, Chou, & Olivia, 2021).

Figure 5 illustrates the classification of demand patterns using the CV^2–p space. The x-axis represents periods between demands (p) in months, a measure of regularity in demand intervals, while the y-axis denotes CV^2, quantifying demand variability. Each point corresponds to a distinct product or SKU, categorized into one of the four demand types based on statistical behavior. The plot reveals the variability and relative frequency of each demand category across the dataset, with Intermittent being the most prevalent (667 instances), followed by Smooth (426), Erratic (270), and Lumpy (233). These classifications provided a critical input for selecting appropriate forecasting models and developing inventory strategies tailored to each demand type.

Figure 5: Material Segmentation

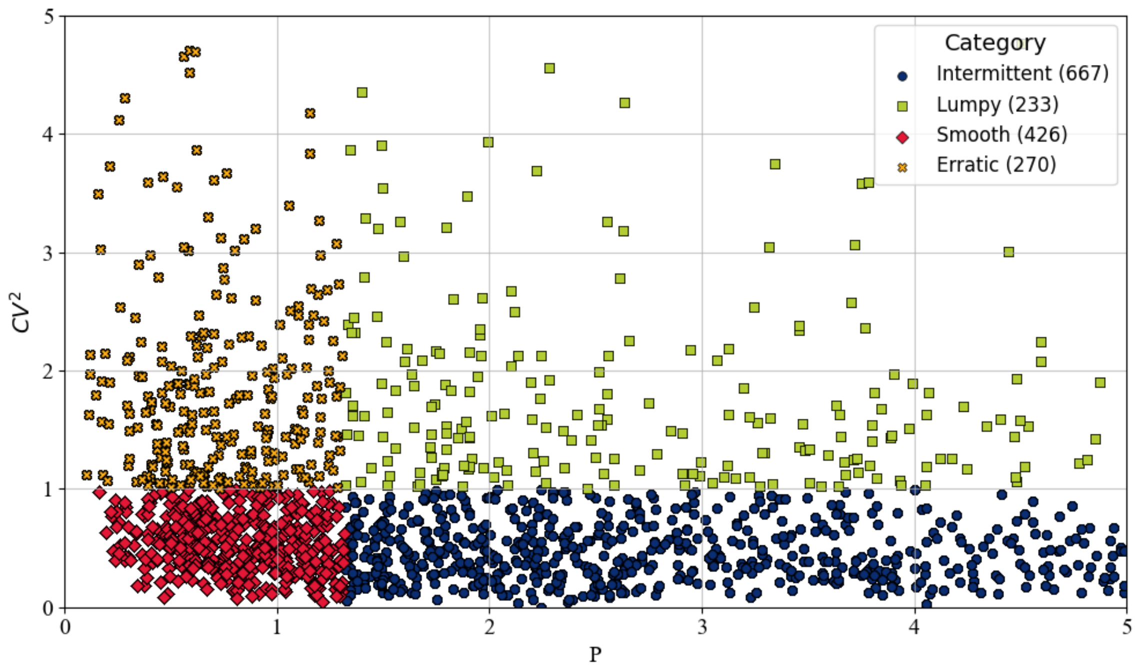

To further illustrate the distinctive characteristics of each demand category, we selected representative materials from each group—Smooth, Intermittent, Erratic, and Lumpy—and plotted their monthly demand patterns over time. As depicted in Figure 6, these time-series plots reveal clear differences in both the frequency of demand occurrences and the variability in demand magnitude. For example, materials classified as Smooth exhibit consistent usage patterns, while Lumpy materials show irregular and highly volatile demand. These examples underscore the importance of aligning forecasting models and inventory policies with the unique behavioral profile of each demand type. A one-size-fits-all approach would fail to capture the operational nuances that drive spare parts consumption, especially in asset-intensive environments such as oil & gas. Tailoring planning strategies based on demand segmentation enables more accurate forecasting and supports more effective inventory optimization.

Figure 6: Representative Time-Series Plots for Each Demand Category

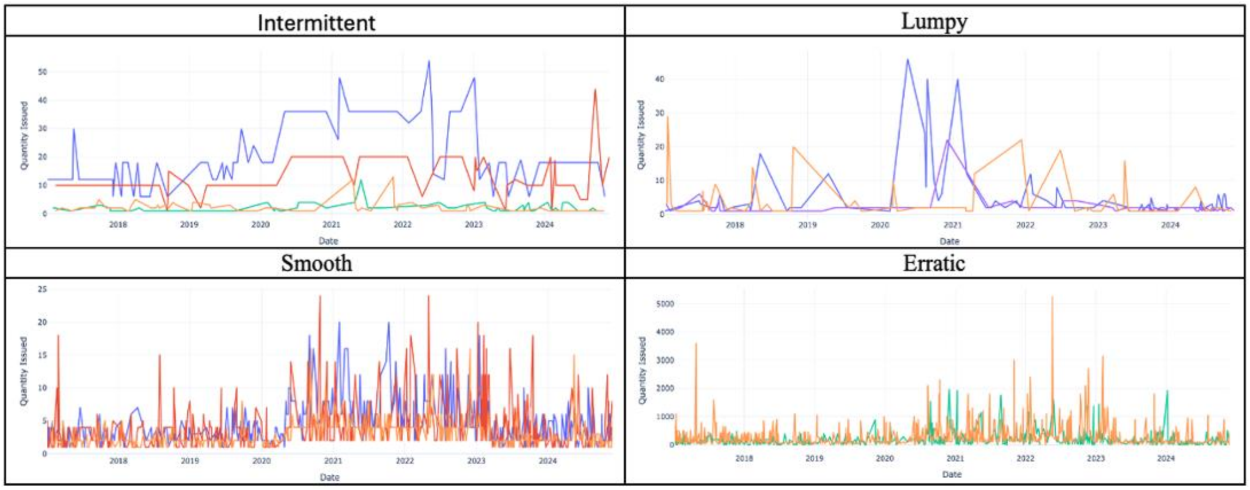

3.5 MACHINE LEARNING – MODEL SELECTION

To accurately forecast spare parts demand and inform appropriate inventory levels for each Stock Keeping Unit (SKU), we employed machine learning (ML) techniques. ML, a key subset of Artificial Intelligence (AI), has become a widely adopted approach for time-series forecasting across numerous domains, including supply chain management. These models use computational algorithms to identify complex patterns in historical data—referred to as training data—and generate predictive outputs for future observations. Model performance is subsequently evaluated using test datasets, enabling objective assessment based on standard error metrics such as Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and R-squared (R^2). In the absence of future or live data, these error metrics serve as proxies to approximate forecast reliability.

Machine learning models can be refined through iterative processes involving hyperparameter tuning, feature engineering, and retraining, thereby improving their predictive robustness. In the context of this project, ML enables a dynamic, data-driven forecasting framework that can account for the variability and intermittence inherent in spare parts demand. Our implementation followed three core steps: (1) model definition and selection, (2) data partitioning into training and testing subsets, and (3) generation and evaluation of monthly demand forecasts for the materials in scope.

3.5.1 DEFINE THE MODEL

According to the No Free Lunch (NFL) theorem, "averaged over all optimization problems, without re-sampling, all optimization algorithms perform equally well" (Wolpert & Macready, 1997). In practical terms, this implies that there is no universal best model, and specialization is required to achieve optimal performance. Applied to spare parts forecasting, we must tailor our models through feature engineering and parameter tuning to ensure that we account for actual demand patterns, stakeholder feedback, and industry specificities.

With the NFL theorem in mind, we recognized that further work was needed to select the optimal ML model. Leveraging the capabilities of the Darts Python library for its unified interface and flexibility (Herzen et al., 2022), we tested various forecasting models, ranging from statistical to deep learning methods. We selected and tuned seven ML models commonly used in demand forecasting, specifically Exponential Smoothing (ETS), AutoARIMA, FourTheta, LightGBM, Croston, Croston SBA, and Prophet. Brief model descriptions below:

- ETS: Captures level, trend, and seasonality parameters to forecast future values based on a weighted average of past observations.
- AutoARIMA: A time series forecasting model that automates the process of identifying the optimal parameters for an ARIMA (Autoregressive Integrated Moving Average) model. It iteratively searches over combinations of p (autoregressive terms), d (differencing terms), and q (moving average terms), as well as seasonal parameters if applicable, to minimize a specified information criterion. By automating parameter selection and differencing, AutoARIMA efficiently models non-stationary time series with both short-term dependencies and long-term trends, making it highly effective for complex, noisy datasets with irregular demand patterns.
- FourTheta: A statistical ensemble method that decomposes time series into trend and residual components by combining multiple Theta models to improve accuracy on complex datasets.
- LightGBM: A gradient boosting framework that models demand using decision trees optimized for speed and efficiency on large datasets.
- Croston's Method: This method is specifically designed for intermittent and erratic demand patterns, where demand occurrences are sparse and have many zero values.
- Croston Syntetos–Boylan Approximation (SBA): A bias-corrected modification of Croston's method that improves forecast accuracy by adjusting for Croston's known positive bias in low-frequency demand scenarios (Syntetos & Boylan, 2005).

- Prophet: A forecasting model that decomposes time series into trend, seasonality, and holiday effects. It handles missing data and outliers effectively, and adapts well to trend changes, making it suitable for time series demand with strong seasonal patterns.

Each model's forecasting performance was assessed against the segmented demand profiles defined in Section 3.4. The goal was to identify the most accurate, robust, and operationally feasible model for each demand segment.

3.5.2 TRAIN AND TEST THE MODEL

Before applying the models, the dataset was split into training and testing subsets. The training set consisted of data captured from January 1, 2022, to December 31, 2023, and the testing set contained the data captured from January 1, 2024, to November 30, 2024. The purpose of this split was to train the models on the historical demand data (training set) to generate a monthly forecast for 2024. The forecasted demand was then compared to the actual known demand (test set), and the residual errors were calculated. Every material was run through all seven models to fit and predict the demand patterns. Figure 7 shows an example of a plot generated for an Intermittent material from the dataset.

Figure 7: ML Performance Assessment for a Material Categorized as "Intermittent"

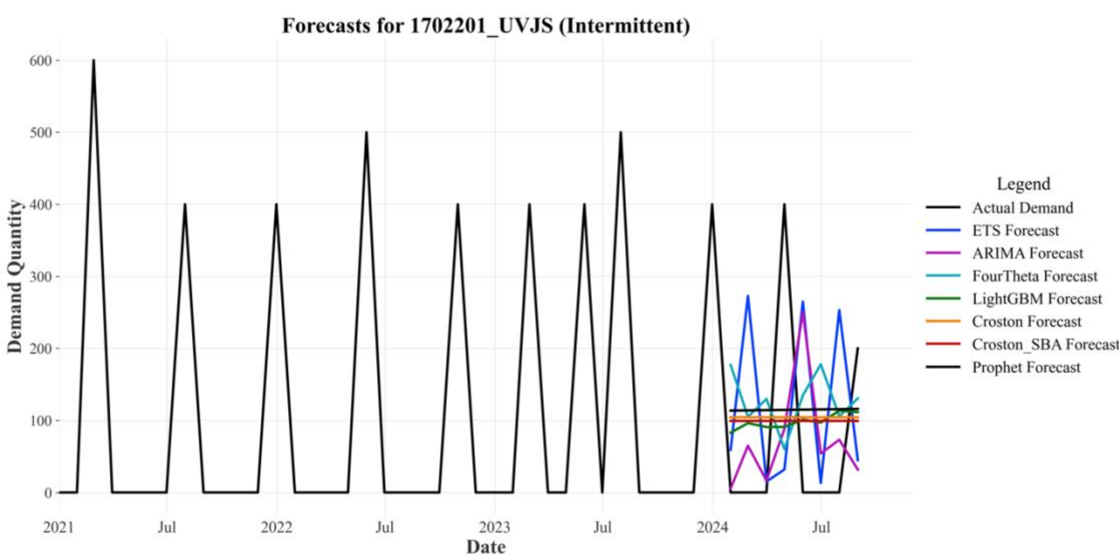

To evaluate the performance and reliability of each forecasting model, we used two widely used error metrics to compare the predicted values against the actual demand:

- Mean Absolute Error (MAE): Measures the average magnitude of errors in the forecast as an absolute value, thus providing a clear indication of accuracy.
- Root Mean Squared Error (RMSE): Emphasizes larger errors by squaring them before averaging, thus providing a more sensitive assessment of the prediction quality.

We developed a ranking model using Mean Absolute Error (MAE), where all forecasting models were ranked from lowest to highest MAE. Table 1 presents the MAE ranking results for the seven forecasting models evaluated across the four distinct demand categories. AutoARIMA outperformed other models in the Intermittent and Lumpy demand categories, exhibiting the lowest MAE value. In contrast, the Croston SBA method achieved the lowest MAE across Erratic and Smooth categories.

Table 1: Comparison of Forecasting Model Performance across Demand Categories

Model	Erratic	Intermittent	Lumpy	Smooth
AutoARIMA	2.94	3.36	2.94	3.54
Croston	3.78	4.19	4.20	3.45
Croston_SBA	2.89	3.62	3.30	3.17
ETS	4.29	4.02	4.23	4.15
FourTheta	3.99	4.21	4.35	3.65
LightGBM	4.99	4.22	4.39	5.57
Prophet	4.77	4.14	4.48	4.34

Given these results, AutoARIMA was selected for Intermittent and Lumpy demand forecasting, and Croston SBA was selected for Erratic and Smooth.

3.6 FORECAST DEMAND

Following comprehensive model assessment and the selection of AutoARIMA and Croston SBA for demand forecasting, we generated 12-month forecasts for the 1,596 materials with sufficient data coverage. These forecasts formed the foundation for developing a dynamic, demand-driven inventory policy tailored to Sponsor's operational needs.

The selected models were run individually for each material in scope, ensuring that the forecasts reflected the specific historical demand behavior for each SKU. Figure 8 presents the forecast example for the same material we assessed in Figure 7.

Figure 8: Demand Forecast Example for a Material Categorized as "Intermittent"

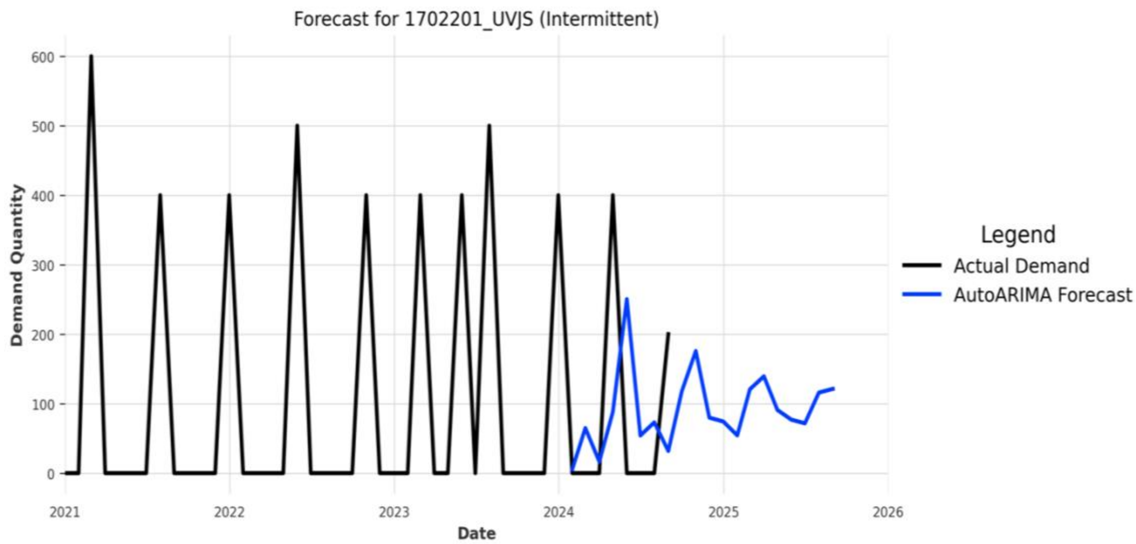

3.7 DEMAND-DRIVEN INVENTORY POLICY

Informed by the work completed for material segmentation, rigid ML model selection, and a 12-month forecast for each of the 1,596 materials in scope, we reached the final project milestone: developing a demand-driven inventory policy designed to enhance service levels and reduce cost.

To establish a robust inventory strategy, we utilized the Root Mean Squared Error (RMSE) values obtained during Demand Forecasting (Section 3.6). RMSE serves as a measure of variability in forecast accuracy, enabling us to compute the appropriate safety stock required for each material. The formula used to calculate safety stock is:

$$Safety\ Stock = k \times RMSE \times \sqrt{Leadtime}$$

Where:

- *k* is the service level factor (e.g., 1.96 for a 95% service level).
- *RMSE* is the forecast error obtained from AutoARIMA and Croston SBA for each material.
- *Lead Time* is average time between purchase requisition and the receipt of purchase order, assumed to be 1 month per Sponsor guidance.

In our calculations, we considered a standard service level of 95%, which corresponds to a k-value of 1.96. However, we recognize that certain materials justify higher service levels based on their criticality. For these materials, the k-values can be easily adjusted – for example, by setting k = 2.33 for 99% or k=2.58 for 99.5%. This flexibility helps reduce the risk of stockouts by aligning inventory policy with material criticality.

Safety Stock is then used to define the recommended inventory per below:

$$Recommended\ Inventory\ Level = Demand\ During\ Lead\ Time + Safety\ Stock$$

After calculating the recommended inventory levels for each material, we compared these results to the Sponsor's current inventory policy, which is based on a fixed three-month demand estimate. This comparison validated the existing policy in certain cases but also revealed widespread understocking risks and overstocking inefficiencies. These findings—and their financial and operational implications—are discussed in detail in Chapter 4: Results and Discussion.

4 RESULTS AND DISCUSSION

4.1 KEY FINDINGS

This Capstone project focused on developing a dynamic, data-driven inventory policy for one of the world's largest energy companies, with the goal of optimizing spare parts inventory, minimizing excess stock, and improving service reliability. Using demand segmentation, machine learning-based forecasting, and flexible inventory recommendations, we analyzed the Sponsor's legacy inventory management approach and compared it to the proposed demand-driven strategy. Based on this analysis, we identified a nearly 27% inventory value reduction opportunity that could unlock working capital to be reinvested into cashflow-generating projects. In parallel, we identified a path to achieving a 95% service level target—enhancing spare parts availability and operational resilience.

Key findings include:

1) Balanced inventory: 22% of all spare parts are adequately stocked, balancing inventory availability and financial investment.
2) Understocking risks: 20% of parts would require an increase in safety stock to ensure sufficient availability. These are the cases where Sponsor's existing inventory did not sufficiently account for demand volume and variability. Insufficient inventory levels for these materials create an elevated risk of stockouts, which could disrupt maintenance activities and impact business continuity.

3) Overstocking inefficiencies: 58% of parts would require a decrease in safety stock to eliminate excess inventory. For these materials, inventory levels significantly exceeded actual operational requirements. The overstocking increased inventory carrying costs, reduced warehouse efficiency, and tied up working capital that could otherwise be deployed for strategic initiatives.
4) Surplus inventory: A notable share of inventory has not moved in over five years. We provided the Sponsor with a detailed surplus analysis to support potential divestment or investment recovery actions.

In summary, the findings reveal a dual opportunity: to improve service levels through targeted inventory increases for critical materials, and to reduce financial waste by divesting or right-sizing non-moving or excess inventory. This approach enables the Sponsor to achieve a projected 27% inventory reduction while maintaining a 95% service level, consistent with industry best practices and organizational targets.

4.2 IMPLICATIONS FOR THE SPONSOR

The findings from this Capstone project have important implications for the Sponsor's operational performance, financial efficiency, and broader strategic objectives.

First, the 27% inventory value reduction opportunity presents a compelling case for unlocking working capital through optimized materials management policies. By divesting surplus inventory and reducing excess holdings while reinforcing availability for critical materials, the Sponsor can redeploy the working capital toward cashflow-generating projects, improving overall financial agility. In an industry characterized by high fixed costs and capital intensity, this reallocation can lead to a strategic advantage.

Second, improving service levels to a 95% target directly supports operational resilience. By increasing safety stock levels for materials identified as understocked, the Sponsor can mitigate the risk of stockouts that may delay maintenance activities, impair asset reliability, and threaten business continuity. Enhancing spare parts availability not only aligns inventory strategy with operational risk management but also strengthens the Sponsor's capacity to maintain high uptime across both upstream and downstream operations.

Third, adopting a dynamic, demand-driven inventory policy represents a shift from static, one-size-fits-all practices to a segmented, data-informed approach. This transition enables a more agile, responsive supply chain capable of adjusting to evolving operational needs and business priorities. By embedding machine learning-based forecasting and inventory segmentation into ongoing planning processes, the Sponsor can build a sustainable capability that improves supply chain maturity over time.

Overall, implementing the recommendations from this project positions the Sponsor to achieve financial benefits, reduce operational risk, and enhance long-term strategic flexibility in spare parts management. The methodology developed in this project is scalable and transferable, with potential applications across other business units, asset classes, or international operations. Additional opportunities for extension are outlined in Section 5.6: Future Research Opportunities.

4.3 LIMITATIONS AND INFLUENCING FACTORS

While this project demonstrates strong potential to improve inventory management through a demand-driven policy, several limitations should be acknowledged.

Data availability and quality constraints influenced model applicability. Of the 40,000 materials in scope, only 1,596 had sufficient historical data to support machine learning-based forecasting. Although these materials represent the most critical, high-volume, and high-spend inventory, the absence of robust historical data for a large proportion of materials limited the scope of advanced modeling. For data-sparse SKUs, average consumption-based safety stock recommendations were applied, introducing simplifications that may not capture all underlying demand drivers.

Forecasting models inherently rely on the assumption that historical demand patterns will persist into the future. While AutoARIMA and SBA models provided strong historical fits, unforeseen operational changes, market disruptions, or shifts in maintenance or turnaround practices could impact forecast accuracy over time. Continuous monitoring and periodic recalibration will be essential to sustain model relevance.

Implementation success is dependent on organizational factors beyond the model itself. Effective application of the inventory optimization methodology will require workflow integration with the Sponsor’s ERP systems, cross-functional adoption, and ongoing governance. Resistance to process changes, lack of data discipline, or inconsistent stakeholder engagement could hinder the full realization of the benefits identified.

Operational realities in the oil & gas sector introduce inherent unpredictability into spare parts demand. Equipment failures, incidents, and unplanned maintenance events can generate demand spikes that are difficult to predict statistically. While safety stock buffers are designed to absorb some of this volatility, no forecasting method can fully eliminate the risks associated with these unexpected events.

Finally, this project focused exclusively on U.S. operations. While the methodology is designed to be scalable and transferable, applying it to international operations with different regulatory, logistical, and operational contexts may require localized adjustments to segmentation thresholds, service level assumptions, lead times considered, and model parameters.

Despite these limitations, the methodology developed provides a scalable, pragmatic, and data-driven foundation for inventory optimization. It equips the Sponsor with a structured framework that can be refined over time as data maturity, system capabilities, and organizational readiness continue to evolve.

5 CONCLUSION

This Capstone project aimed to address a critical supply chain challenge for one of the world's largest energy companies: how to optimize spare parts inventory by aligning stock levels with business demand, minimizing excess inventory, and improving service reliability. By combining domain expertise, advanced analytics, operational insights of our sponsor, and the technical guidance of MIT faculty, we designed and implemented a scalable, demand-driven inventory management framework that is tailored to our Sponsor's operational complexity.

Our work began with comprehensive data preparation, followed by demand segmentation based on demand variability and intermittence metrics. We then tested seven forecasting models across different demand types and selected AutoARIMA for Intermittent and Lumpy segments, while applying the Syntetos–Boylan Approximation (SBA) model for Smooth and Erratic demand. For materials with insufficient historical data, we developed a pragmatic alternative based on a three-month average consumption. This two-pronged approach ensures that inventory recommendations are inclusive, data-driven, and grounded in the operational realities of spare parts management in the energy sector.

The project findings revealed a significant opportunity to reduce inventory investments while simultaneously improving service levels. By transitioning from static inventory policies to dynamic, demand-driven strategies, the Sponsor can unlock working capital, enhance asset readiness, and make operational and strategic decisions that reflect actual business demand patterns and equipment criticality.

5.1 MANAGEMENT RECOMMENDATIONS

Building on the findings of this project, we recommend that the Sponsor organization apply the developed inventory optimization methodology through a phased, scalable approach to drive tangible operational and financial improvements.

First, we recommend implementing the two-pronged inventory policy to ensure material availability and optimal inventory investments. High-volume, data-rich SKUs should be transitioned to the dynamic forecasting models as aligned with the developed methodology, while lower-volume, data-scarce SKUs should adopt the average consumption-based inventory policy. Prioritizing high-volume, high-spend, and critical materials will enable the Sponsor to capture early cost savings and supply assurance benefits. The forecasting model should be run monthly with a Power BI visualization built on top to ensure user friendliness.

Second, we recommend integrating the dynamic safety stock recommendations into the Sponsor's existing ERP system to ensure that inventory levels match the shifts in business demand. This integration will reduce manual intervention, improve reorder point accuracy over time, and facilitate more agile inventory management practices designed to support the ever-changing demands of the business.

Third, we suggest establishing a semi-annual re-evaluation cycle for demand segmentation, model performance monitoring, and safety stock recalibration. Given the inherent volatility of oil and gas operations, maintaining model relevance through continuous improvement practices will be critical to sustaining long-term benefits.

Fourth, we recommend aligning inventory optimization initiatives with cross-functional stakeholders, including operations, maintenance, supply chain, and finance teams. Establishing clear service level targets and ownership for forecast reviews and safety stock updates will help ensure adoption and organizational alignment.

Finally, we recommend piloting the methodology in a single business unit and tracking baseline inventory value, service levels, and working capital impacts before scaling the approach enterprise-wide. A phased rollout will allow the Sponsor to refine change management practices, validate assumptions, and fine-tune parameters in real operational settings.

By adopting this demand-driven inventory policy, the Sponsor can move from a reactive, static planning model to a proactive, dynamic approach that meets business agility requirements and enhances both operational resilience and financial performance.

5.2 FUTURE RESEARCH OPPORTUNITIES

Several areas offer opportunities for further research and development:

- Integration of operational drivers such as maintenance and turnaround schedules, supplier lead times, and capital investment activity into demand forecasting to further improve inventory accuracy.
- Adoption of multi-echelon inventory optimization across the full network of hubs, spokes, upstream and downstream facilities, and the supplier network to capture network-wide efficiencies.
- Integration of additional parts across various scopes, such as operations, maintenance, capital projects, consignment inventory, and third-party managed inventory, to further leverage economies of scale.
- Development of service-level targets for each part in accordance with operational criticality and risk factors.
- Validation and tuning of the methodology for international locations to ensure scalability across diverse operational and regulatory environments.

This project delivers more than just a model: it offers a practical, scalable playbook for transforming spare parts management from a reactive, static process to a proactive, data-driven strategy. By implementing the methodology developed through this capstone, the Sponsor can unlock significant value at scale: reducing inventory by nearly 27%, increasing service levels to support business continuity, and enhancing supply chain agility. Beyond immediate cost savings and service level improvement, this framework helps build supply chain resiliency and positions the Sponsor to adapt to demand volatility, contributing to long-term competitive advantage.

REFERENCES

Bacchetti, A., & Saccani, N. (2012). Spare parts classification and demand forecasting for stock control: Investigating the gap between research and practice. *Omega*, *40*(6), 722–737.

Basten, R. J. I., & van Houtum, G. J. (2014). System-oriented inventory models for spare parts. *Surveys in Operations Research and Management Science*, *19*(1), 34–55.

Boylan, J. E., Syntetos, A. A., & Karakostas, G. C. (2008). Classification for forecasting and stock control: a case study. *Journal of the Operational Research Society*, *59*(4), 473–481.

Chuang, H. H., Chou, Y., & Oliva, R. (2021). Cross-item learning for volatile demand forecasting: An intervention with predictive analytics. *Journal of Operations Management*, *67*(7), 828–852.

Croston, J. D. (1972). Forecasting and Stock Control for Intermittent Demands. *Journal of the Operational Research Society*, *23*(3), 289–303.

Do Rego, J. R., & de Mesquita, M. A. (2015). Demand forecasting and inventory control: A simulation study on automotive spare parts. *International Journal of Production Economics*, *161*, 1–16.

Energy Institute. (2024). *Statistical Review of World Energy*. Statistical Review of World Energy. https://www.energyinst.org/statistical-review.

Fildes, R., Goodwin, P., Lawrence, M., & Nikolopoulos, K. (2009). Effective forecasting and judgmental adjustments: an empirical evaluation and strategies for improvement in supply-chain planning. *International Journal of Forecasting*, *25*(1), 3–23.

Gilbert, J. P., & Finch, B. J. (1985). Maintenance management: Keeping up with production's changing trends and technologies. *Journal of Operations Management*, *6*(1), 1–12.

Grant, D. B., Karagianni, C., & Li, M. (2006). Forecasting and stock obsolescence in whisky production. *International Journal of Logistics Research and Applications*, *9*(3), 319–334.

Herzen, J., Lässig, F., Piazzetta, S. G., Neuer, T., Tafti, L., Raille, G., Pottelbergh, T. V., Pasieka, M., Skrodzki, A., Huguenin, N., Dumonal, M., Kościsz, J., Bader, D., Gusset, F., Benheddi, M., Williamson, C., Kosinski, M., Petrik, M., & Grosch, G. (2022). Darts: User-Friendly Modern Machine Learning for Time Series. *Journal of Machine Learning Research*, *23*(124), 1–6.

IEA. (2024). *Overview and key findings – World Energy Investment 2024 – Analysis*. IEA. https://www.iea.org/reports/world-energy-investment-2024/overview-and-key-findings

Jones, E. C., & Thummalapalli, R. K. (2019). A Methodology to Evaluate Obsolete Inventory in Healthcare. *Supply Chain Engineering and Logistics Handbook, CRC Press EBooks*, 485–513.

Moser, Philipp, et al. "Manufacturing Management in Process Industries: The Impact of Market Conditions and Capital Expenditure on Firm Performance." *IEEE Transactions on Engineering Management*, vol. 68, no. 3, June 2021, pp. 810–822, https://doi.org/10.1109/tem.2019.2914995.

Odedairo, B. O., Alaba, E. H., & Edem, I. (2020). A system dynamics model to determine the value of inventory holding cost. *Journal of Engineering Studies and Research*, *26*(3), 112–123.

Piya, Sujan, et al. “Analysis of Supply Chain Resilience Drivers in Oil and Gas Industries during the COVID-19 Pandemic Using an Integrated Approach.” *Applied Soft Computing*, vol. 121, May 2022, p. 108756, https://doi.org/10.1016/j.asoc.2022.108756. Accessed 27 Apr. 2022.

Rao, A. V. (1973). A Comment on Forecasting and Stock Control for Intermittent Demands. *Journal of the Operational Research Society*, *24*(4), 639–640.

Richardson, H. (1995). Control your costs then cut them. *Transportation and Distribution*, *36*(12), 94–96.

Sanders, N. R., & Manrodt, K. B. (2003). The efficacy of using judgmental versus quantitative forecasting methods in practice. *Omega*, *31*(6), 511–522.

Syntetos, A. A., Babai, M. Z., Y Dallery, & R Teunter. (2008). Periodic control of intermittent demand items: theory and empirical analysis. *Journal of the Operational Research Society*, *60*(5), 611–618.

Syntetos, A. A., Babai, Z., Boylan, J. E., Kolassa, S., & Nikolopoulos, K. (2016). Supply chain forecasting: Theory, practice, their gap and the future. *European Journal of Operational Research*, *252*(1), 1–26.

Syntetos, A. A., & Boylan, J. E. (2005). The accuracy of intermittent demand estimates. *International Journal of Forecasting*, *21*(2), 303–314.

Syntetos, A. A., Boylan, J. E., & Croston, J. D. (2005). On the categorization of demand patterns. *Journal of the Operational Research Society*, *56*(5), 495–503.

Syntetos, A. A., Boylan, J. E., & Disney, S. M. (2009). Forecasting for inventory planning: a 50-year review. *Journal of the Operational Research Society*, *60*(sup1), S149–S160.

Tadayonrad, Y., & Ndiaye, A. B. (2023). A New Key Performance Indicator Model for Demand Forecasting in Inventory Management considering Supply Chain Reliability and Seasonality. *Supply Chain Analytics*, *3*(100026). ScienceDirect.

Thummalapalli, R. (2010). A methodology to evaluate obsolete inventory in health care [Industrial and Management Systems Engineering -- Dissertations and Student Research]. In *DigitalCommons@University of Nebraska - Lincoln*.

Wilkes, H., Schwarzbauer, J. (2010). Hydrocarbons: An Introduction to Structure, Physico-Chemical Properties and Natural Occurrence. *Handbook of Hydrocarbon and Lipid Microbiology*. 1–48.

Williams, T. M. (1984). Stock Control with Sporadic and Slow-Moving Demand. *Journal of the Operational Research Society*, *35*(10), 939–948.

Wolpert, D. H., & Macready, W. G. (1997). No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*, *1*(1), 67–82.