Developing a Dynamic S&OP Process for Third-Party Logistics

by

Richard Augustus Elmquist III
Bachelor of Science in Industrial Engineering, Georgia Institute of Technology, 2008

and

Luis Rodrigo Dávila Novoa
Master of Business Administration, Adolfo Ibañez Business School, 2022

Bachelor of Science in Industrial and Systems Engineering, University of Piura, 2009

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

June 2023

Signature of Author: ____
Department of Supply Chain Management
May 12, 2023

Signature of Author: ____
Department of Supply Chain Management
May 12, 2023

Certified by: ____
Dr. Ilya Jackson
Postdoctoral Associate
Capstone Advisor

Certified by: ____
Dr. Jafar Namdar
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

Developing a Dynamic S&OP Process for Third-Party Logistics

by
Richard Augustus Elmquist III
and
Luis Rodrigo Davila Novoa

Submitted to the Program in Supply Chain Management
on May 12, 2023 in Partial Fulfillment of the
Requirements for the Degree of Master of Applied Science in Supply Chain Management

ABSTRACT

This paper explores the complexities of temperature-sensitive food supply chains and the role of third-party logistics (3PL) providers in managing them. Specifically, we partner with the world's second-largest cold chain 3PL provider to establish a dynamic Sales and Operations Planning (S&OP) process for their warehousing services. In this project, we propose a novel inventory forecasting framework that could complement the S&OP process proposed aimed at aligning supply and demand, optimizing inventory levels, and preserving the quality of temperature-sensitive food. We start by reviewing the related literature on food cold chain management and demand forecasting methods, complemented by qualitative interviews with key roles within the company. We select a specific warehousing site for our minimum viable product and extracted the data of inventory positions for every customer during a 3.5-year period. We propose segmentation criteria based on customer inventory size and variability and develop forecasting models for each segment and used Seasonal Autoregressive Integrated Moving Average (SARIMA) and Facebook Prophet. The accuracy of the models is measured using the Mean Absolute Percentage Error (MAPE) performance metric. The valuable insights offered by the forecasting models allows us to propose create additional freezer capacity at the site. We also identify underutilized space in the cooler segments that could potentially be repurposed to increase freezer capacity. Finally, we discuss next steps regarding the opportunities to improve the forecasting models and scale the dynamic S&OP process across the entire company network. Overall, our findings highlight the importance of a dynamic S&OP process for 3PL providers in managing temperature-sensitive food supply chains effectively. The insights from our inventory forecasting framework can help 3PL providers to optimize their operations, better align supply and demand, and preserve the quality of temperature-sensitive food throughout the supply chain.

Capstone or Thesis Advisor: Dr. Ilya Jackson
Title: Postdoctoral Associate
Capstone or Thesis Co-Advisor: Dr. Jafar Namdar
Title: Postdoctoral Associate

ACKNOWLEDGMENTS

I would like to thank Dr. Ilya Jackson and Dr. Jafar Namdar for their support and motivation throughout the capstone process and overall journey at MIT. I would also like to thank everyone at CTL for their support throughout the master's program. To Rodrigo, my capstone partner for simply being a great teammate and friend throughout the highs and lows of this capstone. To my amazing wife, Natalie, for being my anchor, taking over responsibilities for our two wonderful kids Gus and Pete while I embarked on this journey. To my buddies at MIT, our bond will live forever.

Richard Elmquist

I want to thank Dr. Ilya Jackson and Dr. Jafar Namdar for their friendship and invaluable contribution to our capstone. I would also like to thank the MIT-CTL staff for making us feel part of this great university from the very beginning. To my capstone partner, Richard, for being a great friend and a talented professional. To my wife, Claudia, for making a great sacrifice and making my dreams her own. To my daughters Camila and Rafaela, for being my greatest motivation in life. To my parents, Doris and Hugo, for teaching me the value of education and work. To my brothers Rafael, Daniel, Diego, and Gonzalo for their example. To my mother-in-law, María, for believing in me. Dad, we did it!

Rodrigo Dávila

TABLE OF CONTENTS

1. INTRODUCTION 7
1.1. FOOD INDUSTRY BACKGROUND 8
1.2. COMPANY BACKGROUND 8
1.3. PROBLEM STATEMENT AND RESEARCH QUESTION......9
1.4. GOALS AND EXPECTED OUTCOME......9
2. STATE OF THE ART 11
2.1. FOOD COLD CHAIN MANAGEMENT 11
2.2. SALES AND OPERATIONS PLANNING (S&OP) PROCESS 14
2.3. FORECASTING RATIONALE WITHIN SUPPLY CHAINS......15
2.4. FORECASTING METHODS 16
2.5. TIME SERIES FORECASTING MODELS 16
2.6. EVALUATING FORECAST ACCURACY 19
3. DATA AND METHODOLOGY 21
3.1. DATA COLLECTION 21
3.2. DATA CLEANING......22
3.3. CUSTOMER SEGMENTATION 22
3.4. DATA AGGREGATION 28
3.5. FORECASTING TYPE IDENTIFICATION 34
3.6. FORECASTING METHODOLOGY 39
3.7. BASELINE FORECAST 40
3.8. FORECASTING WITH SARIMA 41
3.9. FORECASTING WITH PROPHET 45
4. RESULTS 50
4.1. FORECAST: ALL CUSTOMERS 51
4.2. FORECAST: COOLER CUSTOMERS 53
4.3. FORECAST: FREEZER CUSTOMERS 54
4.4. FORECAST: CUSTOMER SEGMENTATION 56
5. DISCUSSION 61
5.1. S&OP FRAMEWORK 62
5.2. SCALING ACROSS AMERICOLD NETWORK 63
5.3. INCORPORATING FIXED COMMITMENTS 64
5.4. SELF REGULATED SYSTEMS 65
5.5. FURTHER EXOGENOUS VARIABLES 65
5.6. INBOUND AND OUTBOUND TO FORECAST 66
5.7. ADDING FURTHER CAPACITIES 66
6. CONCLUSION 67
REFERENCES 69
APPENDIX A 70
APPENDIX B......73

LIST OF FIGURES

Figure 1: Structure of Physical Distribution Channels in the Food Sector 14
Figure 2: Measuring Forecast Accuracy 19
Figure 3: Robust Option to Measure Forecast Accuracy 20
Figure 4: Number of Customers per Year 23
Figure 5: Customer Churn Over Time 24
Figure 6: Customer Segmentation Approach 25
Figure 7: Coefficient of Variation (CV) vs. Average Pallets 26
Figure 8: Final Segmentation of Customers 27
Figure 9: Final Customer Segmentation 28
Figure 10: Full Site Inventory Positions 29
Figure 11: Full Site Freezer Inventory Positions 30
Figure 12: Full Site Cooler Inventory Positions 31
Figure 13: Full Site Dry Inventory Positions 32
Figure 14: Aggregated Pallet Positions by Freezer Customer Segments 33
Figure 15: Aggregated Pallet Positions by Cooler Customer Segments 34
Figure 16: Decomposition of Allentown's Inventory time series 35
Figure 17: ACF and PACF of Allentown's Inventory Time Series 36
Figure 18: Correlation Between Lagged CPI and Total Pallet Positions 37
Figure 19: Correlation Between Lagged CPI and Freezer Pallet Positions 38
Figure 20: Correlation Between Lagged CPI and Cooler Pallet Positions 39
Figure 21: Cross Validation: Training vs. Testing Data 40
Figure 22: Fold 1: Cumulative Moving Average 41
Figure 23: Fold 1: SARIMA (1,0,1)(1,1,1)12 43
Figure 24: Residual Analysis: SARIMA (1,0,1)(1,1,1)12 45
Figure 25: Prophet Forecast: Baseline 47
Figure 26: Prophet Forecast: Changepoint = 1 48
Figure 27: Prophet Forecast: Changepoint = 0 49
Figure 28: Prophet Forecast: Growth = Flat 50
Figure 29: All Customers Forecast: SARIMA m(12) (1,0,1) (1,1,1) 52
Figure 30: Cooler Customers: SARIMA m(52) (0,1,0)(1,0,0) 54
Figure 31: Freezer Customers: Prophet – Growth = Flat 55
Figure 32: Cooler Customer Forecast by Segment 57
Figure 33: Consolidated Cooler Customer Segment Forecasts 58
Figure 34: Freezer Customer Forecast by Segment 59
Figure 35: Consolidated Freezer Customer Segment Forecasts 60
Figure 36: Consolidated Full Site Customer Segment Forecasts 61
Figure 37: Proposed S&OP Process 63

LIST OF TABLES

Table 1: Top 5 SARMA AIC Scores 42
Table 2: Top 5 SARIMA MODELS: Full Analysis 44
Table 3: All Customers: Top 3 Models 52
Table 4: Cooler Customers: Top 3 Models 53
Table 5: Freezer Customers Top 3 models 55
Table 6: Top Models: Cooler Customer Segments 56

1. INTRODUCTION

In this segment, we delve into the significance of the food industry and its impact on the US economy. We emphasize the crucial role of the supply chain in this sector, with a particular focus on 3PLs as pivotal players. Furthermore, we introduce Americold, the sponsor company of this capstone, outlining its operations and the challenges it encounters. Lastly, we present the problem statement and research question, which will guide us in achieving the capstone’s expected objectives and outcomes.

1.1. Food Industry Background

The importance of the role of agriculture, food, and related industries in the US is enormous. These industries contributed $1,055 trillion (5%) to the country's GDP in 2020. Moreover, food expenditures of the average household ranked third (11.9% share) in the same year, behind only housing and transportation. Additionally, in 2019, food and beverage manufacturing plants represented 15.8% of the value of shipments of all US manufacturing plants. Meat (24%) and dairy (14%) were their most significant single components (USDA, Food Market and Prices, 2022).

Before ending up on our table, food makes an incredible journey. First, it is harvested and processed by food producers. Then, the food supply chain connects these producers with restaurants, food service providers, and retailers and, in turn, connects the latter with consumers. Managing that supply chain requires considering multiple factors, such as demand, supply, storage, and transportation. In addition, temperature sensitive food needs to be maintained in a temperature-controlled environment to preserve its characteristics and quality. Actors throughout the chain often decide to manage their supply through a third-party logistics (3PL specialized in the refrigerated supply chain. As a result, they leverage the 3PL integrated network and experience, saving time and resources and reducing the risk associated with temperature control (Americold, 2020).

Temperature-controlled products have been growing throughout the last decades. As a sample, the global frozen food market is expected to grow from 314 billion dollars in 2022 to 602 billion dollars in 2032, and the global frozen ready meal market is expected to grow from 40 billion dollars in 2022 to 90 billion dollars in 2032 (Future Marketing Insights. n.d.).

1.2. Company Background

Americold is the world's second-largest refrigerated 3LP provider (2021 IARW Global Top 25 list, Global Cold Chain Alliance, 2021). Their focus is the ownership operation, acquisition, and development of temperature-controlled warehouses. It operates 250 facilities in North America, Europe, Asia-Pacific, and South America, with an overall capacity of approximately 1.5 billion cubic feet and revenue of $2.7 billion (Form 10-K Americold Realty Trust, 2021). The company divides its operations into three segments: Warehouse (storage and handling), third-party managed (management of customer-owned warehouses), and transportation (asset-light consolidation, management, and brokerage services).

Americold serves approximately 4,300 customers. Food manufacturing customers generate 78% of revenue, and retailers generate 20%. Customers come from different sectors, such as retail, packaged foods (17% of revenue), poultry (13%), and dairy (8%). Their 25 largest customers generate 49% of revenue (Americold, Investors presentation, 2022). Moreover, some customer businesses are seasonal, with a gradual increase after May and June (e.g., hot-dog demand in July) and a peak between mid-September and December (e.g., turkey demand in December). This is the reason why physical warehouse occupancy becomes a critical factor in the company´s performance.

Physical warehouse occupancy varies based on several factors, such as throughput maximization and seasonality. To attenuate these variables, Americold seeks to increase fixed commitment in its portfolio, allowing the company to ensure customer space during their peak inventory levels and reduce variability for Americold. Nonetheless, the Sales and Operations departments still face significant challenges trying to maximize economic occupancy without compromising profit and level of service. Some of the challenges include.

1. Anticipating an increase in customer demand and be able to attend to it with current or flexible additional capacity.
2. Maintaining operational productivity when physical capacity is reaching its maximum level.
3. Balancing customer seasonality to smooth occupancy behavior and allocate resources more efficiently.
4. Dealing with occupancy variabilities that require sudden changes in resource allocation (e.g., human resources).
5. Understanding current and future available capacity to allow for new customers that provide incremental revenue for the site.

1.3. Problem Statement and Research Question

There are several factors that can make it difficult for businesses to accurately predict and respond to future demand from their clients. One such challenge is the limited amount of information that clients may provide about their needs and preferences, which can make it hard for businesses to plan and allocate resources effectively. Additionally, the lack of integration between a business's systems and their clients' ERPs can create additional barriers to efficient communication and coordination.

Another challenge is the sudden and unexpected nature of some client demands, which may require businesses to quickly ramp up their capacity or resources to meet these needs. Seasonal variability can also play a role in demand fluctuations, making it difficult for businesses to maintain consistent levels of productivity and output throughout the year. Finally, businesses may struggle to adapt to new changes in demand, such as increasing or decreasing capacity or labor force in the short term, which can result in lost opportunities or reduced profitability. In this context, the overarching question that Americold is trying to answer is:

How can the company maintain operational costs and level of service while attending to current and future demand?

1.4. Goals and Expected Outcomes

Americold chose one site for the scope of this capstone. The selection criterion is their ability to capture the most significant segment of the company's overall network. A site in the Northeast was chosen as the best candidate due to its mixture of customers. This site has a mixture of retail and consumer packaged goods (CPG) customers. It also houses one of the company's largest multi-vendor consolidation (MVC) hubs. A multi-vendor consolidation hub is a centralized site that enables businesses to streamline their supply chain operations by consolidating and managing orders from multiple vendors in one place. The hub acts as an intermediary between the vendors and the retailer, providing a single point of contact for order management, tracking, and communication. The MVC customer uses Americold to consolidate all of their smaller vendors (300) into one location, allowing them to order full truckloads rather than receiving less-than-truckload shipments into their DCs. During the scope development phase, we carefully assessed all potential limitations and identified two distinct types of constraints for distribution centers.

1. **Hard Constraints**: These include inventory locations and dock doors. A dock door is a type of loading dock raised in the air to be level with trucks (Ferrer Commercial Real Estate Advisors. (n.d). *Dock Doors vs Grade Level Doors.* https://ferrercrea.com/dock-doors-vs-grade-level-doors-charleston-commercial-real-estate/). These are considered hard constraints because they have a finite number that cannot change in the short to medium term (less than 6 months).
2. **Soft Constraints**: These include employees and material handling equipment (MHE). These are considered soft constraints as these constraints are not finite. (Ex: Americold could hire more people if the volume continues to grow.) Understanding soft constraints require understanding the impact of productivity on hard constraints.

Upon concluding the scoping phase, we determined that the capstone should primarily center on inventory locations, given that this represents the most challenging constraint to address rapidly, and cannot be readily resolved by implementing additional shifts or overtime.

We hypothesized that using 3 to 5 years of historical data for warehouse inventory on hand can provide a reliable short and medium-term inventory position forecast. Charting these forecasts against known inventory position constraints should provide the company with a view into the future effectiveness of the site.

The company asked that the deliverable include a proof of concept using the above criteria. The proof of concept should include a reliable forecast of inventory. The outcome should give insights into the percent utilization of positions at the site in future time periods. Americold also asked for suggestions on scaling this across all sites within the company based on findings from the proof of concept.

2. STATE OF THE ART

To address the central problem of our capstone – how a Cold Chain 3PL can maintain operational costs and level of service while attending to current and future demand of temperature-controlled warehousing – we reviewed literature in three areas: (1) food cold chain management, (2) Sales and Operations Planning, and (3) demand forecasting methods.

2.1. Food Cold Chain Management

Food assurance is essential for human survival and development (Zhang et al., 2012). To reach the final consumer, food must cross a complex chain that includes farm production, first-line handlers, manufacturers, wholesalers, and retailers. All these actors create the food supply chain. Its management is defined as the "process of getting food products from producers to consumers safely" (Pullman & Wu, 2021).

Food logistics is divided into control activities (transportation oversight, inventory management, production management, and physical network), control execution (transportation and warehousing), and value-added activities (mainly packaging and labeling) (Pullman & Wu, 2021).

As part of the control execution activities, Warehousing allows to manage inventory and accomplishes four purposes: localization (placing processing close to raw materials), decoupling (improving process control), balance (between supply and demand), and buffering (anticipating spikes in demand or delays in production) (Pullman & Wu, 2021).

The food supply chain is facing a range of challenges due to various factors such as changes in countries' economies, evolving lifestyles, climate change, and the impact of pandemics. These challenges have disrupted supply and demand patterns and reduced predictability, necessitating the emergence of new trends to cope with these challenges. One of the significant trends is the adoption of technology, which enhances supply chain transparency and prevents potential issues. Technologies such as data analytics are particularly noteworthy in this regard. Another trend is the shift in customer channels, as online purchasing becomes increasingly popular. Additionally, there is a growing divide between consolidated and localized supply chains. As the food supply chain becomes more specialized, a few multinational companies control most activities (e.g., 80% of sales in each food sector are dominated by four major corporations). However, developing shorter localized food chains can be more agile, responsive, and less prone to contamination (Pullman & Wu, 2021).

A distinguishing feature of the food supply chain is that it must deal with perishability (Zhong et al.; 2016). This means maintaining appropriate temperatures, sanitation conditions, product handling, and avoiding cross-contamination (Pullman & Wu, 2021). Regarding temperature, food must hold proper conditions along the logistics activities to prevent losses (Oliva & Reventria, 2008). In this context, a cold chain is defined as a supply chain that ensures temperature control for perishable products (food, medicines, blood, flowers, etc.). When the product is food, it is known as food cold chain (Sashi et al., 2017).

A temperature-controlled supply chain is "a food supply chain which requires food products to be maintained in a temperature-controlled environment, rather than exposing them to whatever ambient temperatures prevail at the various stages of the supply chain" (Smith & Sparks, 2004, pp.180). Food has four main temperature levels to suit different products: -25 °C for frozen products, +2°C for chilled products, +5°C for produce, and ambient temperature (McKinnon, 2004).

Changes in consumer behavior preferences can explain the growth in temperature-controlled products (e.g., the need to save time and for convenience when acquiring a frozen meal). Also, changes in supply systems and actors enabled these preferences to be attended (e.g., technological capabilities that allow products to be brought across the world to satisfy out-of-season demand) (Smith & Sparks, 2004).

Many food companies outsource their supply chain activities to third-party logistics (3PL) operators. 3PL is defined as "the use of external companies to perform logistics functions that have traditionally been performed within an organization" (Lieb, 1992). Those functions involve freight transport, warehousing, inventory management, and materials handling individually or in various combinations (McKinnon, 2004). The United States Census Bureau dimensions the importance of this sector in its 2017 report. In that year, more than 200,000 transportation and warehousing employer businesses and more than 100,000 self-employed truck drivers generated 909 billion dollars of revenue and 5 million employments.

As shown in Figure 1, Food supply chain can be divided into three levels where 3PL can participate by providing services: primary, secondary, and tertiary. At the primary level, it can provide trunk haulage, primary consolidation, and pallet-load services. At the secondary level, it can participate in integrated contract distribution, transport and warehousing services, and reverse logistics. At the tertiary level it can supply food deliveries to independent retailers and catering outlets (McKinnon, 2004).

Figure 1

Structure of Physical Distribution Channels in the Food Sector

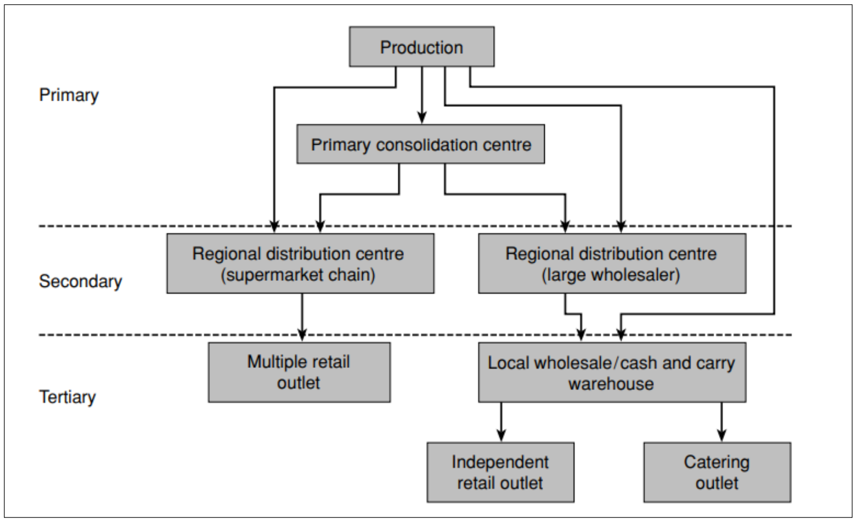

Note. Structure of physical distribution channels in the food sector. Reprinted from *The Food Supply Chain* (p. 168), by A.C. McKinnon, 2004, Blackwell Publishing Ltd. Copyright 2004.

The decision to outsource logistical activities is mainly taken from a cost perspective. 3PL operators can reduce costs by consolidating and balancing fluctuations in clients' traffic, achieving a higher traffic utilization level. Other benefits are avoiding asset investments, reducing risks related to their ownership, and upgrading the standard of operations by leveraging the 3PL operator experience. This last point is particularly important in the cold food supply chain because of the perishability and fragility of products, time-sensitivity of deliveries and the variability in throughput related to seasonality (McKinnon, 2004)

2.2. Sales and Operations Planning (S&OP) Process

Sales and operations planning (S&OP) is a key business process that evaluates expected customer demand against supply capabilities of the organization. This process provides a framework for coordination between different functions of the organization. Most research in this field focuses on companies that manufacture or procure materials. The overall goal of S&OP is to match customer demand with the ability for the company to manufacture or procure the desired goods (Tuomikangas & Kaipia, 2014).

Americold, like many Third-party logistics (3PL) companies, focus on warehousing. In Americold's case this removes two key areas of the supply chain, manufacturing and point of sale. This focus reduces the visibility to data that typically drives an S&OP process, and thus hinders the ability to create a robust S&OP process. In recent years, research has shown that a closer integration with supply chain partners (3PL) can positively impact a firm's performance. A significant portion of these studies look at information integration. Although information integration has shown to positively affect both the focal firm and its partners, companies are mostly underutilizing supply chain integration in practice (Jayaram & Tan, 2010).

For this capstone we will focus on using internal 3PL data to drive a robust S&OP process. This will require a need to forecast warehouse inventory to understand customer capacity needs. We will overlay these forecasts with site capacity constraints at the site to understand the likelihood of being able to execute to this demand.

2.3. Forecasting Rationale within Supply Chains

A supply chain encompasses all relevant parties that satisfy demand of the final customer. A "party" can be any corporation or entity that makes decisions on behalf of the customer. The ordering pattern of the final customer creates demand signals upstream through the supply chain. In a generalized supply chain, retailers send demand signals upstream to wholesalers and/or distributers which, in turn, place orders to manufacturers. These manufacturers create demand signals further upstream to procure raw materials and/or components (Syntetos et al., 2016). If final customer demand were consistent or known well in advance, a supply chain would be a basic scheduling exercise. Customer demand, in practice, is uncertain, which drives the need for forecasting across the supply chain.

2.4. Forecasting Methods

Nearly all decisions made within corporations consider some type of forecast. The three general methods of forecasting are: qualitative techniques, time series analysis, and causal methods (Chambers et al., 1971).

- **Qualitative analysis**: Used when little information is known. Techniques such as questioning experts (Delphi Technique) or market research are used to create an approximate forecast for the future.
- **Time series analysis**: Focuses exclusively on historical data. The goal of this technique is to find patterns within the data that can be projected on the future.
- **Causal methods**: Blend historical data with specific information that has a clear relationship with the system being forecasted. Using gross domestic product (GDP) numbers to help forecast retail sales would be an example of a causal relationship. The goal is to find information that is highly correlated with the company's' data being forecasted.

For the remainder of this section, we will focus on time series forecasting models. Specific forecasting models will allow us to add causal factors within the model using exogenous factors. We will also discuss qualitative analysis as a way to create a more robust S&OP process as subject matter experts could have site specific changes not accounted for in the historical data.

2.5. Time Series Forecasting Models

Time series forecasting models have been a major area of study over the last few decades. These models use four key components that provide the basis for the forecast: level, trend, seasonality, and random fluctuations. (MITx_MicroMasters_SCM_KeyConcepts.Pdf, n.d.)

- **Level:** Characterized by a constant pattern over time. With no other pattern present, a level forecast would be a straight line over time. Commonly referred to as stationary.

- **Trend:** Characterized by a constant rate of growth or decline over a given time horizon. A trend can be linear or non-linear.
- **Seasonality:** Characterized by a repeated pattern over a fixed time horizon. These patterns can show up daily, weekly, monthly, yearly.
- **Random fluctuations:** Account for the remainder of the variability in a model. These fluctuations are considered "noise" or "errors" and form the basis for forecast accuracy.

Among the most popular forecasting models are exponential smoothing models. These models were proposed in the late 1950's and have been the basis for many successful models. Exponential smoothing models' weigh new observations more in the overall forecast with older observations decaying exponentially over time. This decay is created by using a smoothing parameter for level, trend, and seasonality. For the trend parameter there is an underlying assumption that the trend stays constant over time. In practice this is not always true, especially for long horizon forecasts. To combat this issue a dampening parameter can be added to the equation (Hyndman & Athanasopoulos, 2018). For our capstone we expect to use the Holt-Winters model, which accounts for trend and seasonality. We will also validate if a dampening parameter provides a better forecast.

Autoregressive integrated moving average (ARIMA) models provide another path to forecast time series data. These models were developed in the 1970's by George Box and Gwilym Jenkins. While exponential models only focus on describing the trend and seasonality within the data, ARIMA models also focus on describing the autocorrelations within the data. The three properties of an ARIMA model are p, d, and q and can be denoted as ARIMA(p,d,q). The first property p is the autoregressive factor on the time series and can be denoted as AR(p). Autoregression models forecast the variable using a linear combination of previous data points, essentially regressing the variable against itself. The value p refers to the number of periods used in the ARIMA model. The second property d is the differencing factor on the time series. Most time series have some form of trend and seasonality meaning they are non-stationary. Differencing turns a non-stationary time series into a level or stationary time series. The value of d refers to the number of transformations needed to create a stationary time series. The third variable q is the moving average factor and can be denoted as MA(q). Moving average models use a weighted moving average of past forecast errors. This allows the forecast to account for errors in the past. The value q refers to the number of periods of forecast error used in the ARIMA model (Hyndman & Athanasopoulos, 2018).

It is important to note that the ARMIA model discussed above assumes there is no seasonality within the data. In our capstone we expect seasonality to be present in the data. To account for seasonality, we will use the Seasonal ARIMA model or SARIMA. It is written as seen in (1).

$$ARIMA \quad (p,d,q) \quad (P,D,Q)_m \tag{1}$$

The variable *m* pertains to the number of seasons observed within a year. The uppercase variables account for the seasonal part of the model while the lowercase variables account for the non-seasonal part of the model (Hyndman & Athanasopoulos, 2018).

One study also found that using exogenous variables within the SARIMA model (SARIMAX) created a better forecast for bananas within the food retail sector (Arunraj et al., 2016). In this study a SARIMA model was able to account for only 38.6% of the variation in demand. Using holidays, price reductions, and months as exogenous factors, the SARIMAX model was able to account for 61.3% of the variation. In the study, all exogenous variables had a positive effect on the model with the largest being price reductions. The sites chosen for this capstone will not have point of sale data, so price reduction data will not be easily accessible. The study still provides useful insight in the ability to use holidays and months as options for this capstone. Understanding casual relationships, we could use exogenous variables to help create a better forecast.

In recent years, new methods have been created that seek to simplify forecasting. One of the most important is Facebook's Prophet model. A few major advantages of Prophet are accommodating seasonality across multiple periods, no removal of outliers, model fitting is fast, and the model has easily interpretable parameters. The basic formula is as seen in (2).

$$y(t) = g(t) + s(t) + h(t) + e_t \tag{2}$$

Where g(t) models' trend for non-periodic changes, s(t) represents periodic changes (e.g. weekly and yearly seasonality), and h(t) measures the effects of holidays which are normally irregular. The error term accounts for all variances in the model. A saturating growth model and piecewise linear model were used as the basis for the trend model within Prophet (Taylor & Letham ,2018). The saturated growth model is appealing for this capstone as the variables being forecasted will have a fixed capacity. Prophet also allows for exogenous variables.

2.6. Evaluating Forecast Accuracy

It is important to measure forecast accuracy based on actual forecasts. Available data is commonly split into two portions, training and testing data. Training data is used to estimate parameters of the forecast model and test data is used to measure the accuracy of the forecasting model. Typically, 80% of the data is used for training and 20% of the data is used for testing (Hyndman & Athanasopoulos, 2018).

Figure 2

Measuring Forecast Accuracy

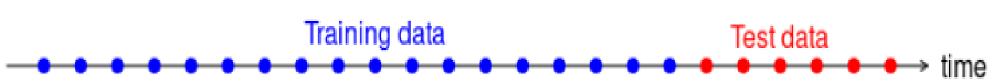

Note: Provides example of training versus testing data to compute forecast accuracy. From "*Forecasting: principles and practice*, 2nd edition", by Hyndman & Athanasopoulos, 2018.

A more complex version of training/test data is time series cross-validation. In this process there are a series of training sets that consist of only one forecast. The preceding data can only be utilized in generating the present forecast.

Figure 3

Robust Option to Measure Forecast Accuracy

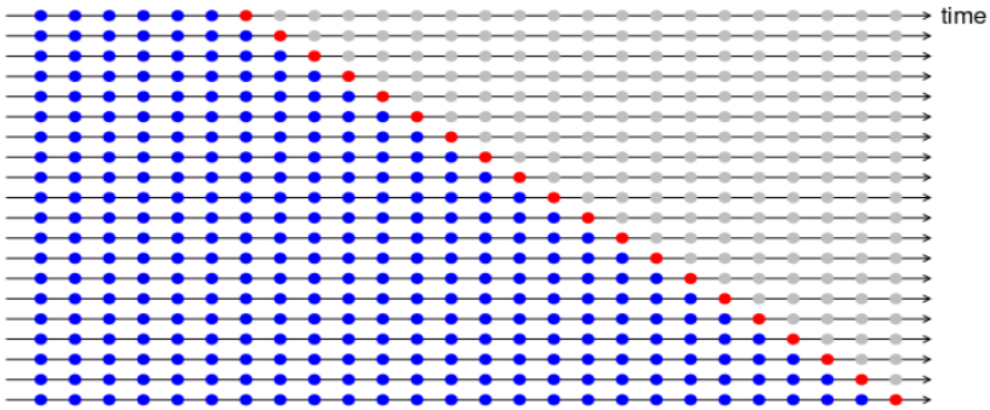

Note: Provides more robust example of training versus testing data to measure forecast accuracy. From "*Forecasting: principles and practice*, 2nd edition", by Hyndman & Athanasopoulos, 2018.

A forecast error is calculated by taking the difference between the forecast and the actual data observed in the test data. Forecast errors can be summarized as a single unit using scale-dependent errors or percentage errors. Scale dependent errors use the same scale as the underlying data. Percentage errors are useful because of their ability to be unit free and are beneficial when looking at forecast accuracy across different units (Hyndman & Athanasopoulos, 2018). In this capstone, both scale-dependent errors such as Mean Absolute Error (MAE) and Root Mean Square Error (RMSE) as well as percentage errors like Mean Absolute Percentage Error (MAPE) will be employed. The formulas for each are written as in (3).

$$Mean\ absolute\ error: MAE = mean(|A_t - F_t|)$$

$$Mean\ absolute\ percent\ error: mean(\left|\frac{A_t - F_t}{A_t}\right|)$$

$$Root\ mean\ squared\ error: RMSE = \sqrt{mean((A_t - F_t)^2)} \tag{3}$$

$$A_t = Actual\ forecast\ at\ time\ t$$

$$F_t = Forecast\ at\ time\ t$$

As per Hyndman & Athanasopoulos, (2018),choosing the best forecast that minimizes the MAE or MAPE tends to result in forecasts that are closer to the median, while using the RMSE leads to forecasts closer to the mean. A crucial characteristic of a good forecast is to have residuals that are uncorrelated with a mean of 0. Correlated residuals may contain additional information that could enhance the forecast accuracy, while a non-zero mean of residuals may indicate a biased forecast. In our capstone, we aim to determine the forecast that minimizes the MAPE while also ensuring uncorrelated residuals with near zero bias. This will be crucial in achieving accurate and reliable forecasts for capacity issues over the forecast time horizon.

3. DATA AND METHODOLOGY

This section describes the steps we took to collect and clean the data to be ready for processing. In addition, due to the large volume of data and its variability, we applied segmenation criteria that allowed us to analyze each group more homogeneously. Finally, we described the different forecasting methodologies used and the reasons why they were appropriate for our dataset**.**

3.1. Data Collection

Americold provided data from 4 distinct tables within their database. To ensure scalability in the future, all columns and rows for each table discussed below was provided for the desired timeframes. Due to size constraints, multiple CSV files were created for each distinct table. The first phase of the capstone was to consolidate like files into a local database. The local database consisted of the five tables shown below:

- **Outbound Demand Data**: Four years (2019 -2022) of data consisting of line level sales order history. Key fields within this data are customer id, Item number, mode of transportation, and quantity shipped.
- **Inbound Supply Data**: Four years (2019 – 2022) of data consisting of line level inbound orders. Key fields within this data are customer id, Item number, mode of transportation, and quantity received.

- **Warehouse Inventory Data**: Forty-two months of data (July 2019 – December 2022) of aggregated data. The data was aggregated due to the size of the files. Data was aggregated by day based on customer and item.
- **Customer Master**: This file provides key customer data information that allows us to create a standard unit of measure across all data sets. The standard unit of measure to be used in our forecasting models will be pallets as pallets are the standard unit of measure moved throughout the facility.

3.2. Data Cleansing

Over the course of this Capstone, multiple tests were conducted on the data to ensure its accuracy and reliability. As a result of this data validation process, we identified two specific areas that required further refinement of the raw data.

Firstly, we discovered the presence of duplicate data across all files for April 21, 2021, which would have impacted the accuracy of our findings had they not been addressed. Consequently, we promptly removed the duplicates for this day, thus ensuring the integrity of our subsequent analyses. We also added a function into the model to check for this type of issue in the future.

Secondly, we encountered issues with approximately 20% of items in the item master file that was provided, which in turn caused difficulties in accurately comprehending pallet quantities. Despite this setback, however, we were able to successfully address the issue by examining other sites within the Americold network that catered to similar customers. Through this process, we were able to obtain the correct item master records.

3.3. Customer Segmentation

After setting up the local database and conducting necessary data cleaning procedures, we were able to identify 385 individual customers who possessed stock at the site during the period covered by the dataset. Further analysis of the unique customers per year indicated that the maximum number of customers within a single year was 276, as illustrated in Figure 4. This posed a significant challenge since generating 385 distinct customer forecasts would result in significant forecasting inaccuracies and render the outputs incomprehensible.

Figure 4

Number of Customers per Year

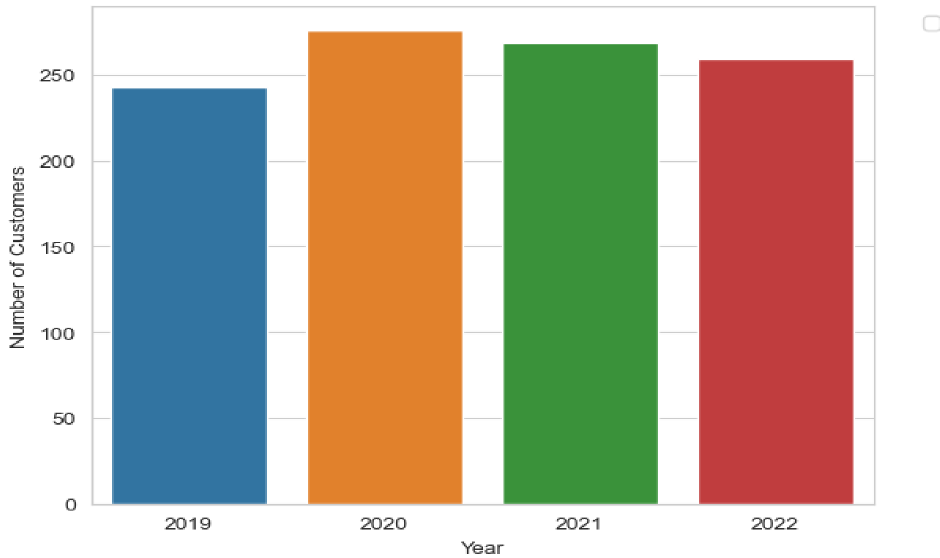

Note: Bar chart provides total number of customers present within a given year over the 4-year time horizon

Subsequently, we examined customer churn at the site by categorizing customers into two groups each year. We referred to a customer as an "Existing Customer" if they had been present the preceding year, while customers without inventory on hand in the previous year were labeled as "New Customers." As depicted in Figure 5, new customers were present every year, except in 2019 when we lacked data for the preceding year. In 2020, 73 new customers had inventory at the site, accounting for 26% of the total 276 customers with inventory. The presence of new customers every year further complicated the forecasting model due to the reduced historical data available for forecasting.

Figure 5

Customer Churn Over Time

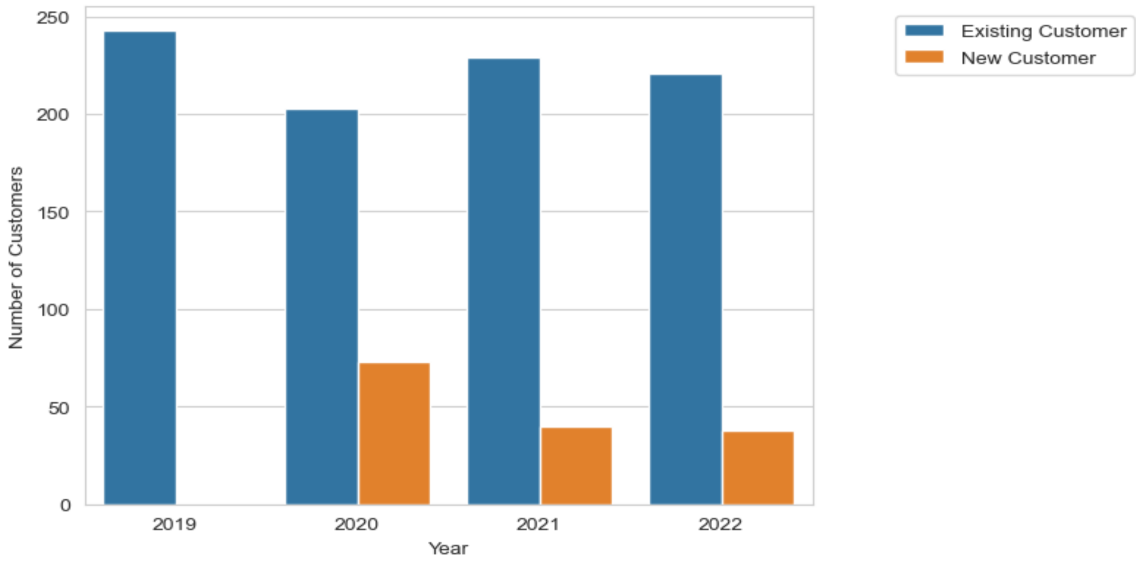

Note. Bar chart showing new customers vs. existing customer. A new customer is defined as a customer that was not present the previous year. All customers are assumed to be "Existing Customers" in 2019 as no data was available for 2018.

Segmenting was deemed crucial for generating useful forecasts for the organization due to the presence of new customers annually and overall number of customers. We explored several alternatives and eventually developed our segmentation model inspired by the Boston Consulting Group's (BCG) Growth Share Model. The BCG model is based on a two-by-two matrix with market share and market growth as its primary factors for segmentation. In our model, we replaced these factors with inventory size and ease of business as our two main criteria. (See Figure 6) To determine inventory size, we computed the average inventory during the analysis year. To identify the analysis year, we used the current date when the code was executed. If the date was more than 6 months into the year, the current year was selected; otherwise, the previous year was chosen as the analysis year. This approach ensures scalability for future analyses. For ease of business, we utilized the Coefficient of Variation (CV) for outbound demand. The CV is calculated by dividing the standard deviation by the mean, which indicates the dispersion around the mean. The lower the dispersion, the more dependable the volume is on any given day. We selected outbound demand for this metric as it is the most labor-intensive activity in a distribution center.

Figure 6

Customer Segmentation Approach

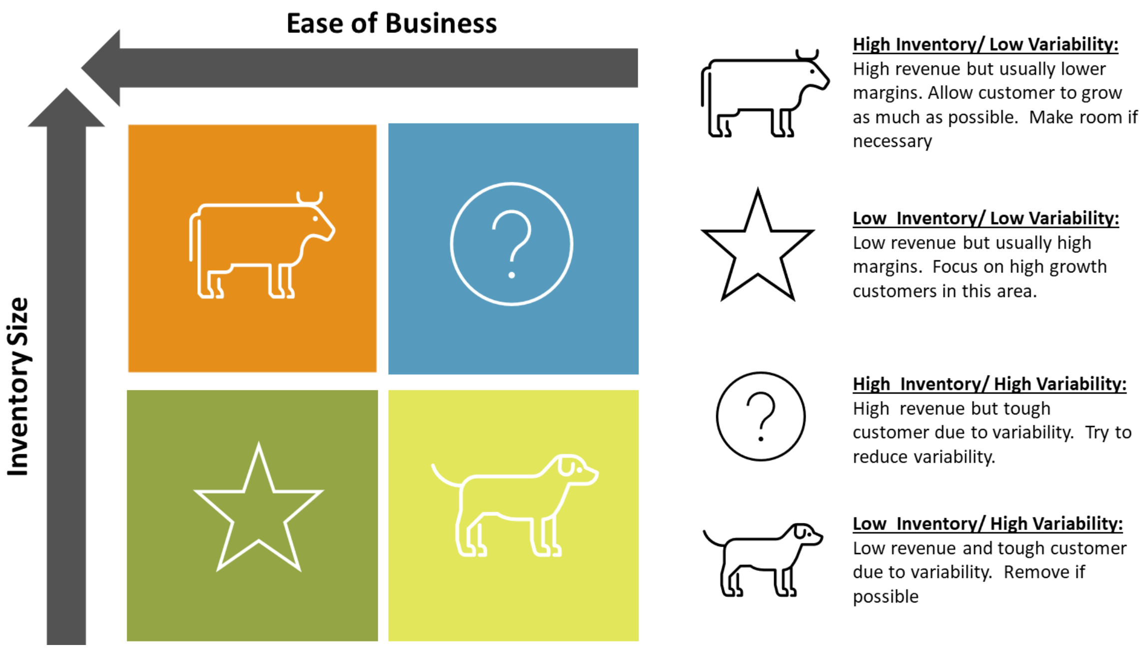

Note. Graphic was created internally but inspired by the BCG growth share matrix to help segment customers

Using this approach with the data provided by Americold, we generated the output shown in figure 7, where each point represents all 385 customers in the data set. The Y-axis displays the average pallets for the analysis year, which was identified as 2022 because the current date was less than 6 months into 2023. In 2022, the largest customer had an average of 5,521 pallets on hand. The X-axis shows the Coefficient of Variation for each customer in the dataset, calculated across the entire data set, while the average inventory was calculated only for 2022. This approach was intentionally adopted to consolidate customers no longer present at the site.

Figure 7

Coefficient of Variation (CV) vs. Average Pallets

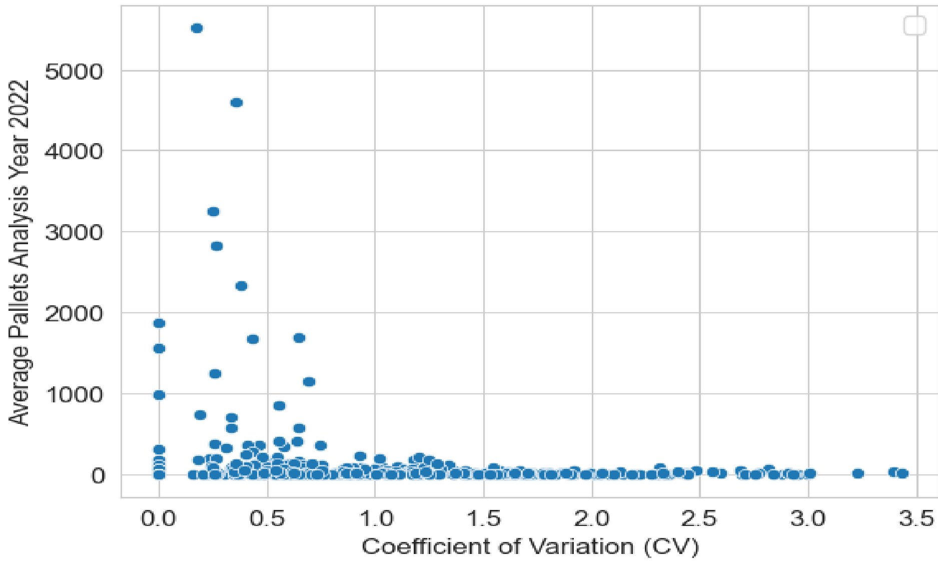

Note. Scatterplot showing average pallets in 2022 (Y axis) vs. the Coefficient of Variation (X axis) for each customer within the data set.

Upon completing a comprehensive data analysis, we found that certain customers exhibited a high Coefficient of Variation (CV), rendering it unfeasible to divide them into quadrants of equal size based solely on their CV values. To address this issue, we created a function that could calculate the minimum value between the maximum CV divided by 2 or 0.5 in order to facilitate scaling. The formula for is written as in (4)

$$CV\ Split = Min\ (\frac{\max(CV)}{2}\ , 0.5) \tag{4}$$

We chose a CV value of 0.5 as the maximum split because, in the context of forecasting, a lower degree of dispersion generally indicates a more stable and predictable data pattern. This, in turn, can lead to more accurate and reliable forecasts. Our primary objective with selecting a CV of 0.5 as the maximum split was to ensure that we had robust forecasts for the "Low CV" customer segments.

Regarding pallet positions, we used the maximum pallet positions divided by 2 to create evenly sized quadrants on the Y-axis. To provide a visual representation of the final segmentation created by our model, Figure 8 was produced.

Figure 8

Final Segmentation of Customers

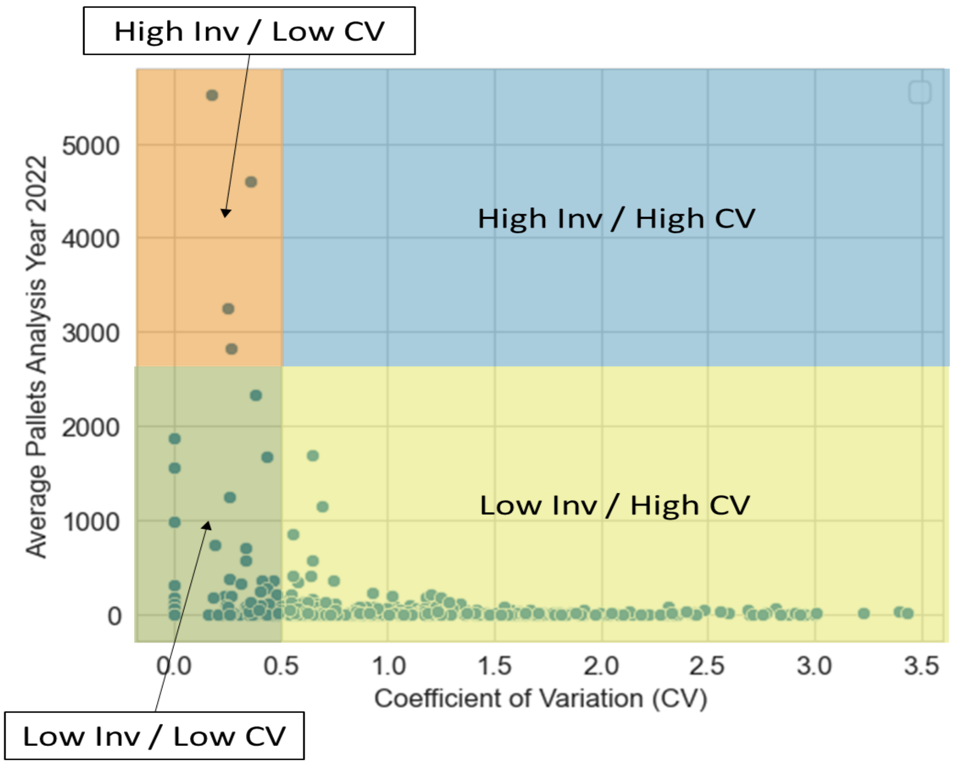

Note. Scatterplot showing average pallets in 2022 (Y axis) vs. the Coefficient of Variation (X axis) for each customer within the data set overlayed with the customer segments created. (Low Inv/ Low CV, High Inv / Low CV, Low Inv / High CV, High Inv / High CV)

To finalize the segmentation for each customer, three additional features were required for proper scaling of the model. Firstly, using the Item Master table, we identified the storage temperature for each customer item which allowed us to categorize all items into Freezer, Cooler or Dry. Secondly, we identified customers within the multi-vendor consolidation program through a column in the Item Master. Regardless of their category, customers were given an MVC program flag along with the temperature stored. Lastly, we created unique customer IDs (e.g. key customer 1) for all customers in the high inventory and low CV group. Americold believed it was important to track these customers individually due to their significance to the site. With these features, we were able to reduce the total number of unique customers from 385 to 11 actionable segmentation groups using this new customer segmentation model (See figure 9).

Figure 9

Final Customer Segmentation

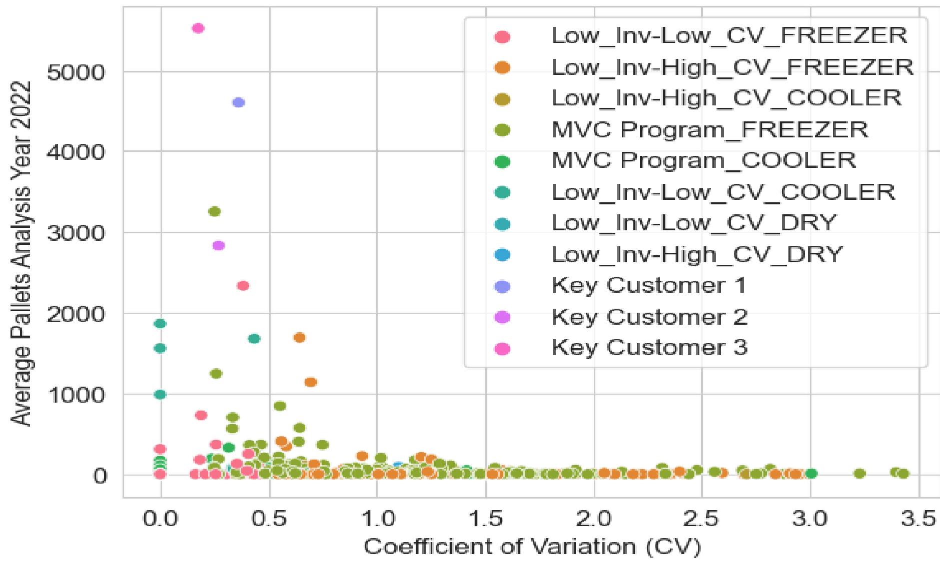

Note. Scatterplot showing average pallets in 2022 (Y axis) vs. the Coefficient of Variation (X axis) for each customer within the data set with each customer color coded based on their customer segments.

3.4. Data Aggregation

In our capstone, inventory management was our top priority since it was the most critical constraint to track on-site. To start data exploration and visualization, we needed to determine the data frequency to use. We had daily inventory data up to that point. During our discussions with the company, we found that Americold required a 26-week visibility in advance to make medium-term decisions. After several rounds of aggregation, we concluded that weekly aggregation was the most effective approach, where we consolidated the daily data into a weekly mean inventory data for each customer segment. Daily data was too erratic for our forecasting models, while monthly aggregation removed critical peaks necessary for modeling inflexible constraints that could not be altered in the short term. We identified three useful aggregations for forecasting to address the customer's problem statement of maintaining operational costs and service levels while meeting current and future demand.

The initial valuable aggregation involved examining the inventory positions throughout the entire site. In Figure 10, the aggregated inventory on hand across the site is depicted throughout the dataset. By aggregating data across the entire site, the company can obtain a broad overview of inventory positions across the customer base, enabling Americold to keep track of site inventory in relation to inventory capacity limitations. While this approach may be helpful, it does not offer the level of detail required for this site as the site has multiple temperature zones.

Figure 10

Full Site Inventory Positions

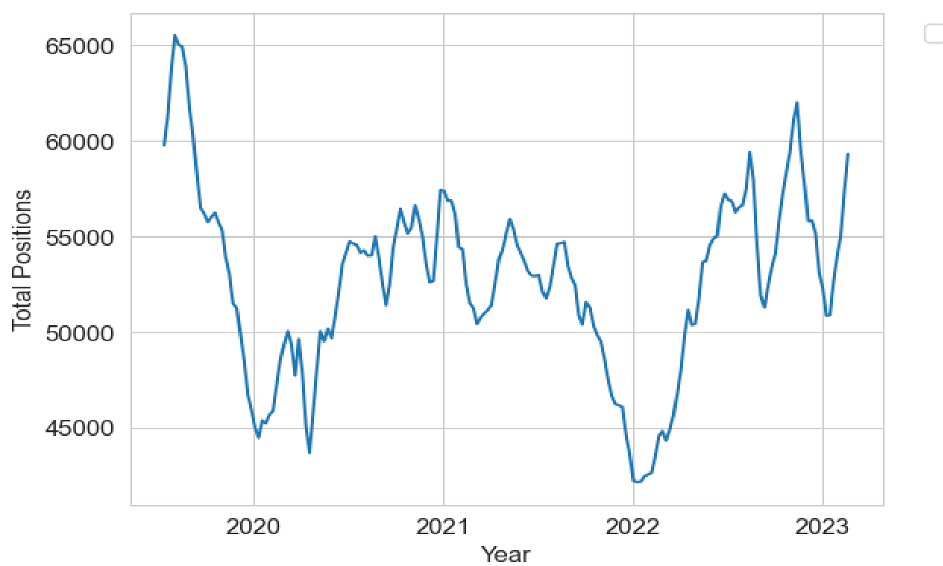

Note. Aggregated total pallet positions for all customers within the site.

After breaking down the total site inventory into different temperature ranges, the subsequent level of aggregation was achieved. The site comprises three distinct temperature zones: Freezer, Cooler, and Dry (Ambient), which cannot be stored together in the same room, thereby creating additional constraints. By analyzing this lower-tier data, we can identify whether there is a need to adjust the temperature within a room to cater to the growth of a specific temperature range.

The aggregated inventory on hand for freezer customers across the site throughout the dataset is depicted in Figure 11. The majority of the inventory falls under the freezer temperature range, which is a significant driver of changes in overall site inventory.

Figure 11

Full Site Freezer Inventory Positions

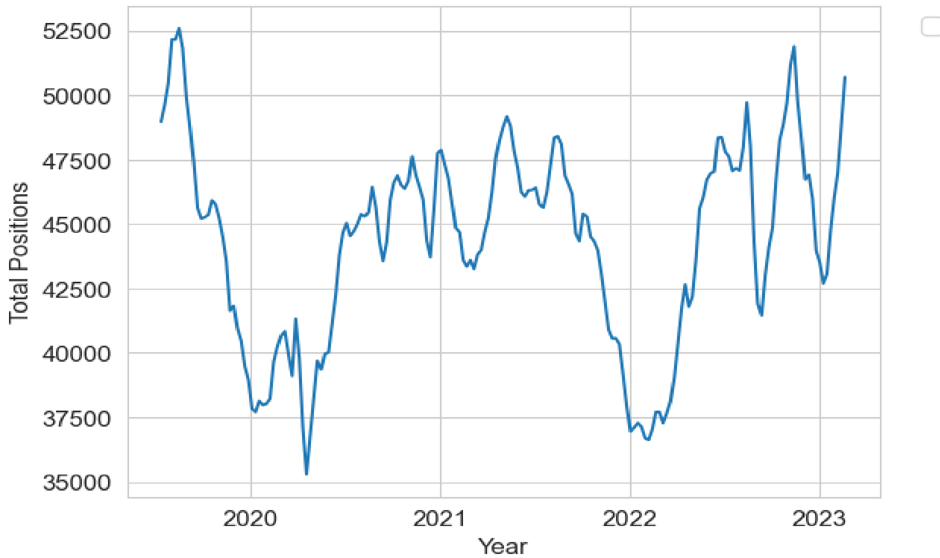

Note. Aggregated total pallet positions for all freezer customers within the site.

The aggregated inventory on hand for cooler customers across the site throughout the dataset is depicted in Figure 13. Cooler on average occupies around 16% of the actual site pallet positions.

Figure 12

Full Site Cooler Inventory Positions

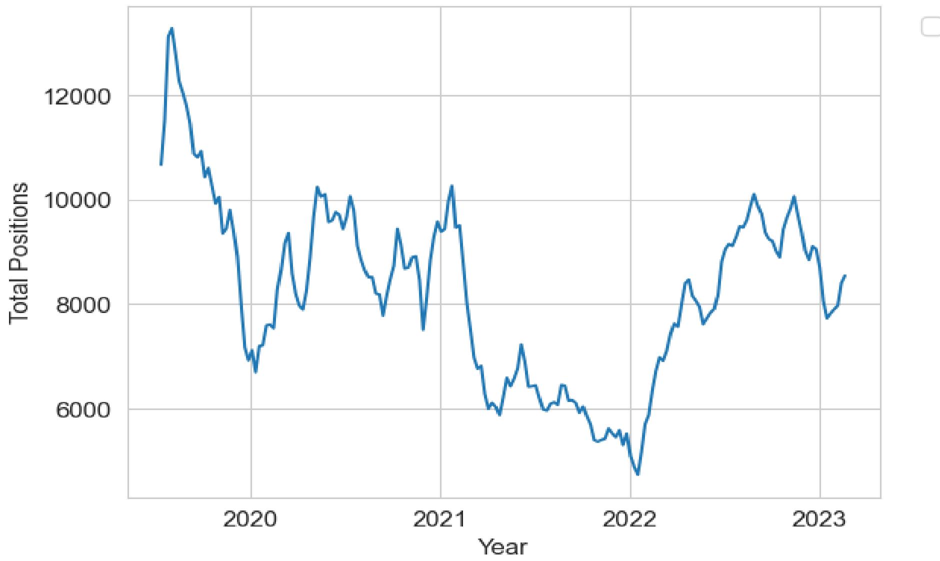

Note. Aggregated total pallet positions for all cooler customers within the site.

The aggregated inventory on hand for Dry (ambient) customers across the site throughout the dataset is depicted in Figure 13. Dry (ambient) customers account for the smallest portion of pallet positions. Americold does not prioritize ambient customers strategically, and these pallet positions were made available primarily because existing customers required a small amount of ambient space. As of 03/01/2023, there are no more available Dry (ambient) pallet positions, and there are no plans to create more in the future. Hence, we will not be conducting any further analysis on this category. Despite the absence of Dry (ambient) pallet positions, we will still include overall pallet positions in our analysis since they did occupy space in the site. Additionally, we will maintain the functions that were created in the overall model for the purpose of scaling.

Figure 13

Full Site Dry (Ambient) Inventory Positions

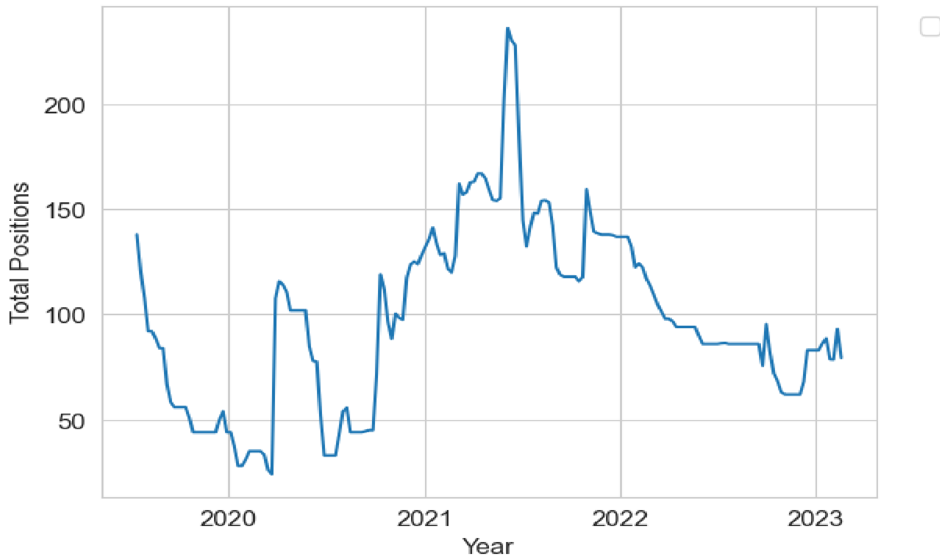

Note. Aggregated total pallet positions for all dry customers within the site.

The final step of aggregation involved grouping the data by customer segments established in section 3.3. It is crucial for the company to understand the customer characteristics that impact inventory levels, whether positively or negatively. For example, if a key customer (with high inventory and low CV) is experiencing rapid growth, Americold would need to analyze this customer differently than a collection of customers with low inventory and high CV. This level of granularity enables the company to make informed strategic decisions, such as expanding the facility to accommodate further growth. To better comprehend the 9 (2 segments associated with dry customers) customer segmentations, figures 14 and 15 are presented by temperature range.

The freezer portion of the building consists of six customer segmentations, as illustrated in figure 14. The MVC program at the site (MVC_program_FREEZER) has increased their pallet positions by almost 10,000, while all other customer segments remain relatively unchanged or decreased. In section 1.3, we discussed Americold's strategy of securing fixed pallet commitments from customers, which enables better management of inventory fluctuations and guarantees customers space during peak periods. Although fixed commitments are out of scope for this capstone, the presence of stable inventory suggests that Americold may be artificially limiting customer growth. Fixed commitments will be further explored in section 5.3.

Figure 14

Aggregated Pallet Positions by Freezer Customer Segments

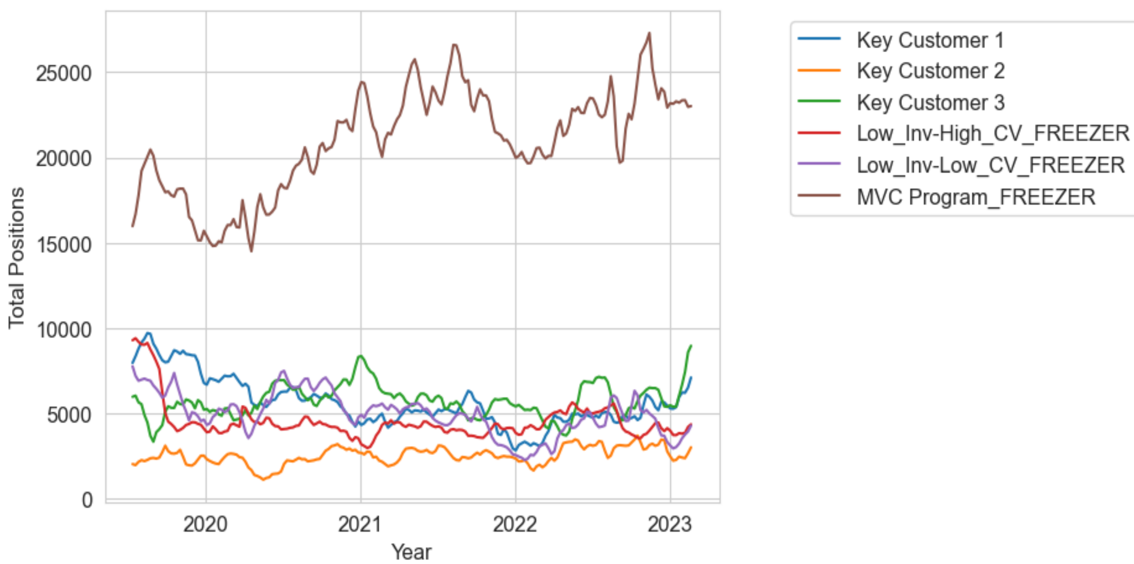

Note. Aggregated total pallet positions for all segments associated with the freezer business within the site.

The cooler portion of the building consists of three customer segmentations, as illustrated in figure 15. The cooler MVC program at the site (MVC_program_Cooler) also saw an increase over this timeframe. Starting from a lower base of around 1K the customer segment pallet positions doubled. The Low inventory and low CV customer segment (Low_INV-Low_CV_COOLER) saw a significant decrease from 2019 to the beginning of 2022, before increasing to around 2/3 of the peak in 2019. We were informed by the site that one of the cooler rooms was converted into a freezer due to this decline.

Figure 15

Aggregated Pallet Positions by Cooler Customer Segments

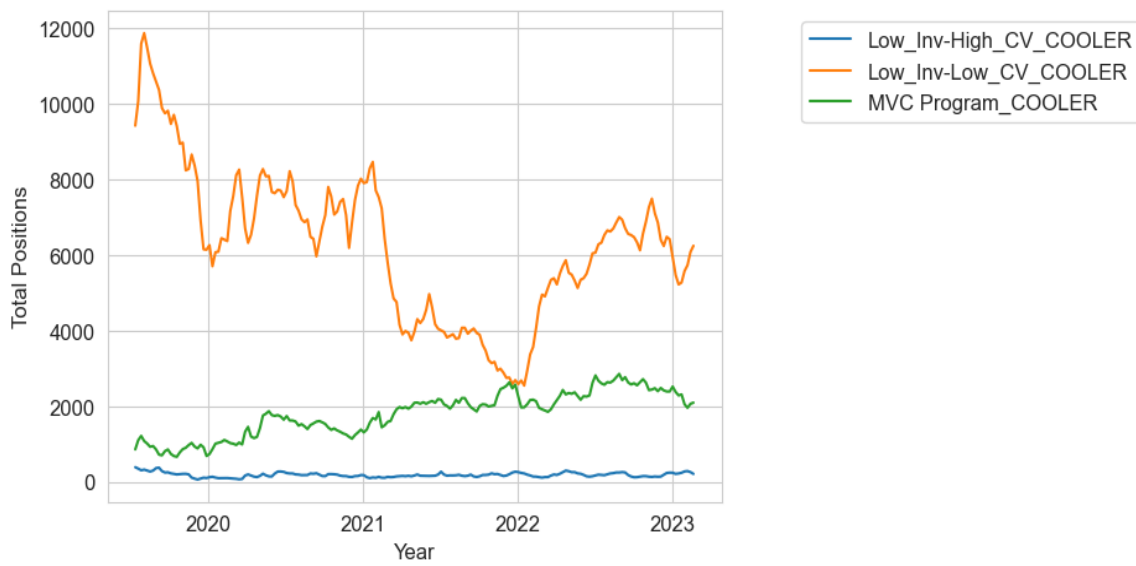

Note. Aggregated total pallet positions for all customer segments associated with the cooler business within the site.

3.5. Forecast Type Identification

Our next step was to identify the appropriate forecasting methods to use in our analysis. After analyzing the full site inventory in Figure 10 to gain general insights, we utilized decomposition analysis in Figure 17 to better understand the time series characteristics. During the analysis period, the values in the full site inventory varied from 42,122 to 65,521 pallet positions, with a mean of 52,618 pallet positions and a standard deviation of 4813. As seen in figure 16, we divided the time series into trend, seasonality, and residuals to gain more information about the data. The trend component was not uniform throughout the time series, with positive slopes from May 2020 to January 2021 and from January 2022 to August 2022, and negative slopes from January 2020 to April 2020 and February 2021 to December 2021. Nonetheless, we noticed a cyclical pattern in the trend's behavior. Using 52 periods within the year, we found that the seasonality component could account for an additional 4,000 pallets in September to a reduction of 4,000 pallets in February. Residuals were dispersed around zero and were unbiased, indicating good quality. They ranged from 6,000 pallet positions to -6,000 pallet positions.

Figure 16

Decomposition of Allentown's Inventory Time Series

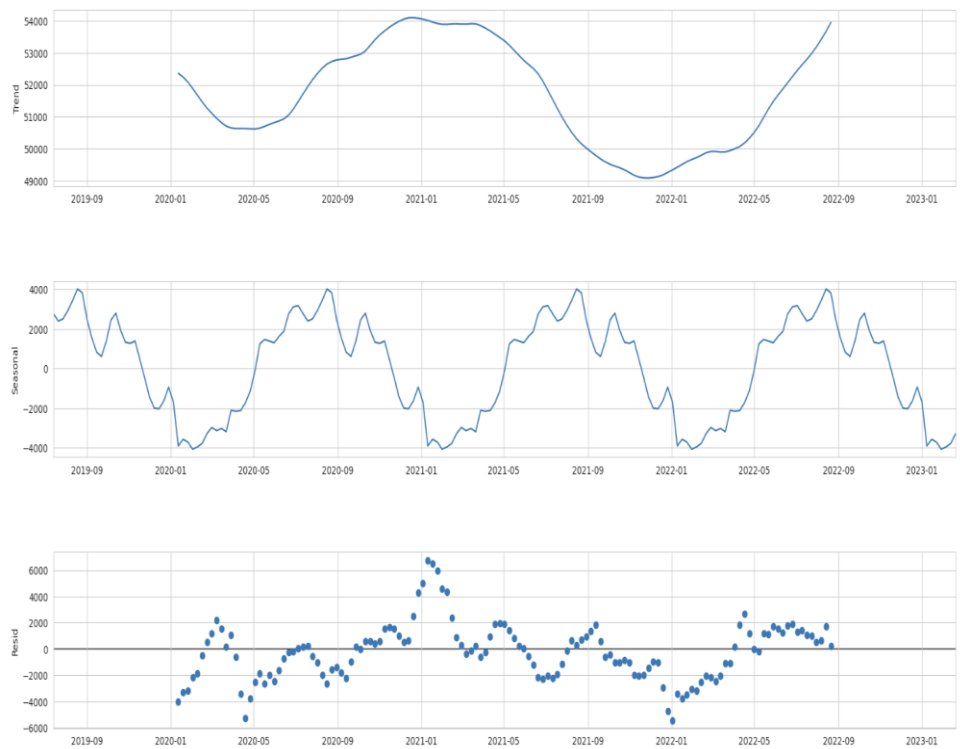

Note. Decomposition analysis for aggregated total pallet positions over all customers within the site.

We next conducted an Autocorrelation Function (ACF) and a Partial Autocorrelation Function (PACF) to identify potential correlations between time series values and their lags. As depicted in Figure 17, the results showed positive autocorrelation with the first seven lagged values, positive partial autocorrelation with the first lagged value, and negative partial autocorrelation with the second lagged value.

Figure 17

ACF and PACF of Allentown's Inventory Time Series

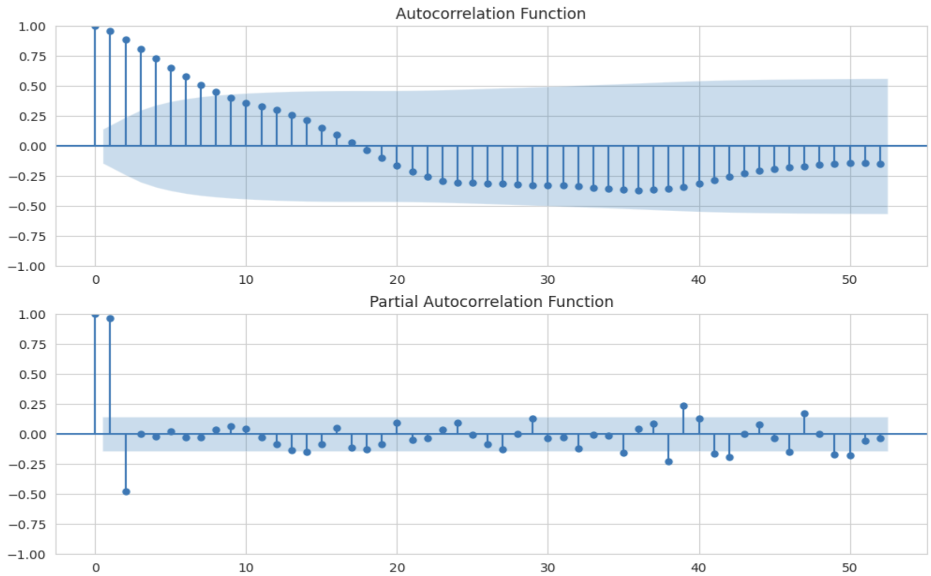

Note. ACF and PACF plots for aggregated total pallet positions over all customers within the site.

Our last step to understand proper forecasting models was to identify any exogenous factors that could help the forecast. We conducted a search for external (exogenous) factors to determine their potential correlation with inventory positions. Correlation coefficients measure the strength and direction of a linear relationship between two variables and range from -1 to +1. A coefficient of -1 indicates a perfect negative correlation, where one variable decreases as the other increases; a coefficient of +1 indicates a perfect positive correlation, where both variables increase simultaneously; a coefficient of 0 suggests no correlation, meaning there is no linear association between the two variables. However, a correlation coefficient of 0 does not necessarily imply that there is no relationship between the variables, only that the relationship is not linear. We identified five exogenous factors that we believed could be significant: Consumer Price Index (CPI), Gross Domestic Product (GDP), total COVID Cases, COVID Cases above 1MM, and COVID lockdowns. To analyze these factors, we created 25 lags for each variable, corresponding to 25 weeks. Lags are a common technique in data analysis and modeling, especially in time series analysis, where data points are observed at regular intervals over time. A lagged variable is a variable that has been shifted back in time by a certain number of time periods. We used lags because changes in one variable might not immediately impact the target variable, inventory positions in this case.

To illustrate this concept, consider a hypothetical example. Suppose that in 2020, there was an increase in the number of COVID cases, resulting in the closure of customer A's manufacturing plants. However, this closure did not have an immediate effect on inbound volume, and as a result, inventory levels were impacted only after several weeks. In such a scenario, using lags would enable us to determine the optimal number of weeks it takes for these plant closures to affect pallet positions.

The correlation between lagged CPI and total pallet positions is illustrated in Figure 18. The x-axis indicates the number of lagged weeks, while the y-axis shows the corresponding correlation coefficients. The plot reveals low correlation coefficients, suggesting a weak or non-existent correlation between these variables.

Figure 18

Correlation Between Lagged CPI and Total Pallet Positions

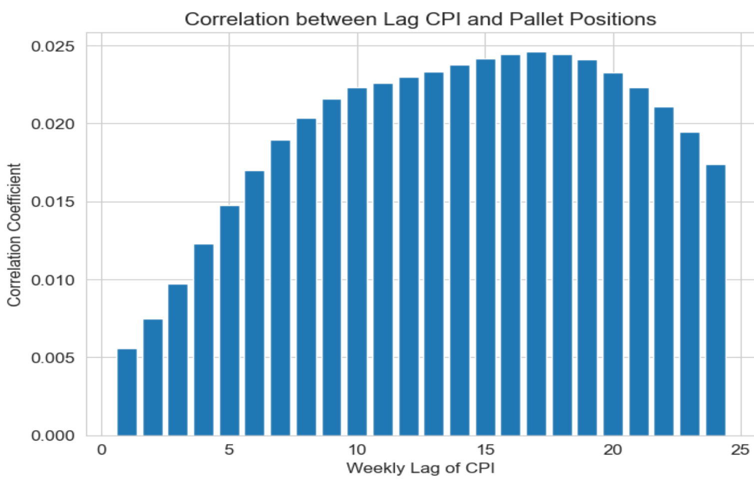

Note. Correlation between lagged CPI and total pallet positions for all customers at the site is represented in a correlation matrix. The X-axis shows the number of weeks CPI was lagged, while the Y-axis displays the percentage correlation for those lags.

The correlation between lagged CPI and Freezer pallet positions is illustrated in Figure 19. Similar to Figure 18, the graph shows low correlation coefficients, indicating a weak or non-existent relationship between these two variables.

Figure 19

Correlation Between Lagged CPI and Freezer Pallet Positions

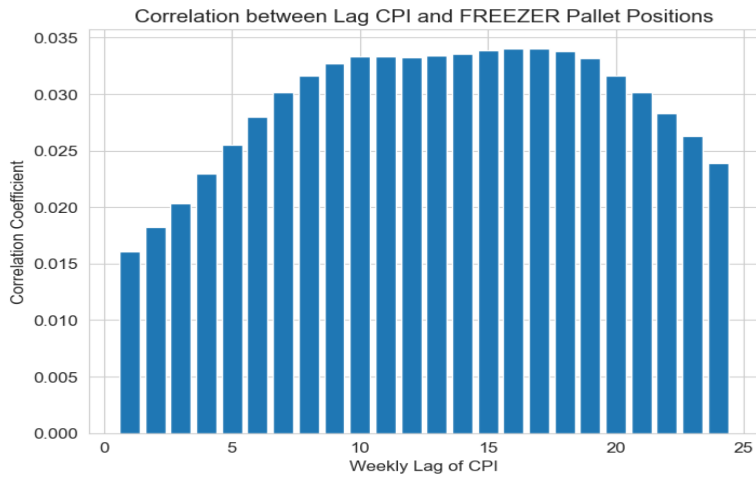

Note. Correlation between lagged CPI and total pallet positions for all freezer customers at the site is represented in a correlation matrix. The X-axis shows the number of weeks CPI was lagged, while the Y-axis displays the percentage correlation for those lags.

The correlation between lagged CPI and cooler pallet positions is illustrated in Figure 20. The high correlation coefficients in the plot suggest that there is no relationship between these variables.

Figure 20

Correlation Between Lagged CPI and Cooler Pallet Positions

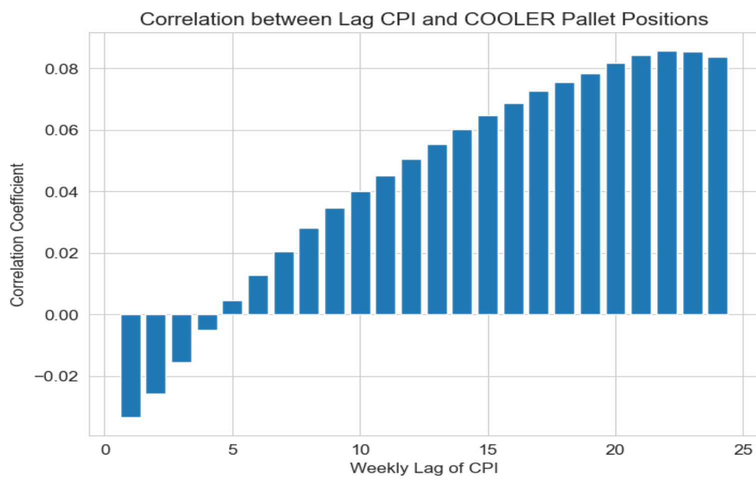

Note. Correlation between lagged CPI and total pallet positions for all cooler customers at the site is represented in a correlation matrix. The X-axis shows the number of weeks CPI was lagged, while the Y-axis displays the percentage correlation for those lags.

We conducted a similar analysis on the remaining identified variables: Gross Domestic Product (GDP), total COVID Cases, COVID Cases above 1MM, and COVID lockdowns. However, our findings indicated that these variables had minimal to no correlation with the inventory position. In other words, we did not observe a significant linear relationship between any of these variables and the inventory position.

After performing these analyses, we were able to select suitable forecasting methods. SARIMA, which is a seasonal autoregressive integrated moving average model, was chosen due to the presence of seasonality and autocorrelations. Additionally, we opted to use Facebook's Prophet, which also considers seasonal factors and autocorrelations. Given the longer forecast horizon required for this capstone, we opted not to use Holt-Winters.

3.6. Forecasting Methodology

The methodology adopted in the State of the Art chapter aimed to identify the best forecasting model by minimizing the mean absolute percent error (MAPE), while also considering the root mean square error (RMSE) and mean absolute error (MAE) of all forecasts presented. We used a fourfold cross-validation technique, ensuring that the first fold contained over a year of data to account for seasonal factors. Figure 21 provides a visual representation of the folds used in all models, with the x-axis indicating the number of weeks of data.

Figure 21

Cross Validation: Training vs. Testing Data

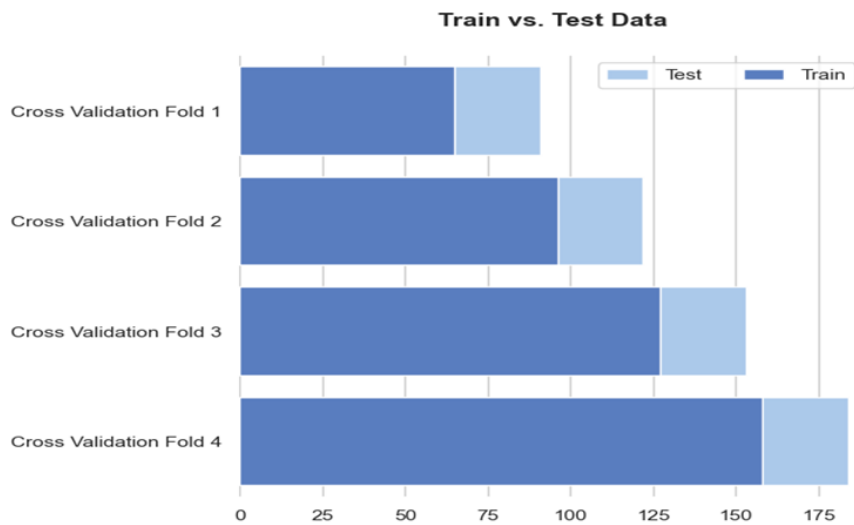

Note. Graph showing cross validation structure used across all models. The x-axis corresponds to the number of weeks for training and testing.

Each test cycle comprised 26 weeks, and we calculated the MAPE over this period, averaging the MAPEs for each fold to determine the final model accuracy. In sections 3.7, 3.8, and 3.9, we will discuss the parameter selection process for the SARIMA and Prophet models, focusing on the total site inventory aggregated at the week level. We will also examine our baseline forecast, which serves as a reference point for evaluating the performance of our other models.

3.7. Baseline Forecast

In order to evaluate the performance of our models, we required a baseline model, which we established as a cumulative moving average using the same cross-validation technique outlined above. For the 26-week testing period, we maintained the cumulative moving average at a constant level.

Figure 22 displays the initial cross-validation fold of our cumulative moving average model, representing the total site inventory at the weekly level. The corresponding mean absolute percentage error (MAPE) is also included in the illustration. The mean MAPE across all folds was 6.94% for total site inventory aggregated at the week level.

Figure 22

Fold 1: Cumulative Moving Average

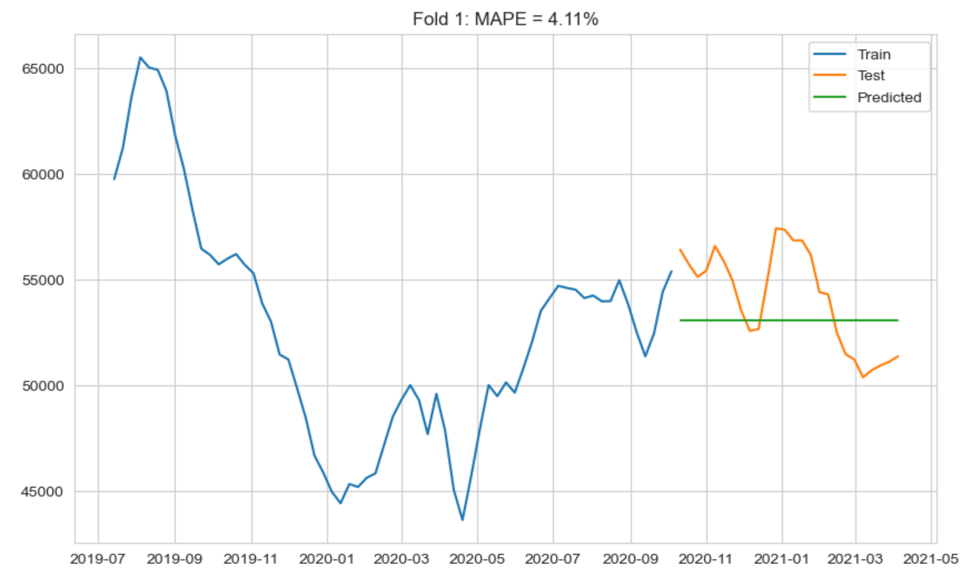

Note. Initial cross-validation fold used to evaluate the accuracy of the cumulative model. The actual data is represented by the orange line, while the forecast is depicted by the green line for total pallet positions across all customers.

3.8. Forecasting with SARIMA

To make predictions using SARIMA models, we must provide various input parameters such as non-seasonal parameters (p, d, q) and seasonal parameters (P, D, Q), where m represents the number of time steps in a seasonal period. Our study explored two options for m: m = 12 for a 12-week seasonal period and m = 52 for a one-year seasonal period to account for quarterly and yearly seasonality.

The next step was to determine the remaining parameters. We chose multiple combinations for p, d, q, P, D, Q. Each parameter took an integer value from 0 to 2 in the case of SARIMA with m=12 and from 0 to 1 in the case of SARIMA with m=52.

This created 252 possible combinations for our SARIMA models for a specific data aggregation. To select the best model, we fitted each combination and performed an Akaike Information Criteria (AIC) test to measure the goodness of fit and model simplicity. We chose the top 50% combinations with the lowest AIC score for each model type (SARIMA with m=12 and SARIMA with m=52). Table 1 provides the top 5 parameter combinations sorted by lowest AIC score for SARIMA models with m=12.

Table 1

Top 5 SARMA AIC Scores

Combination	p	d	q	P	D	Q	m	AIC
1	2	0	0	2	1	0	12	2,567
2	1	0	0	2	1	1	12	2,624
3	1	0	2	1	1	1	12	2,763
4	2	0	0	1	1	0	12	2,777
5	1	0	1	1	1	1	12	2,784

Note. Table providing top 5 SARIMA models sorted by best (smallest) AIC score for total pallet positions across all customers.

We utilized time series cross-validation and the MAPE metric to assess the performance of the selected models. The model with the lowest MAPE was chosen for each type. In addition, we examined the coefficients displayed by the model and checked their significance by verifying if their p-values were less than 0.05. If any coefficient had a p-value greater than 0.05, we discarded the model and selected the one with the second lowest MAPE. We repeated this process until we found a model where all coefficients were significant. Figure 23 illustrates the first cross-validation fold for the best parameters for SARIMA models with m=12 baseline and its corresponding MAPE.

Figure 23

Fold 1: SARIMA(1,0,1)(1,1,1)12

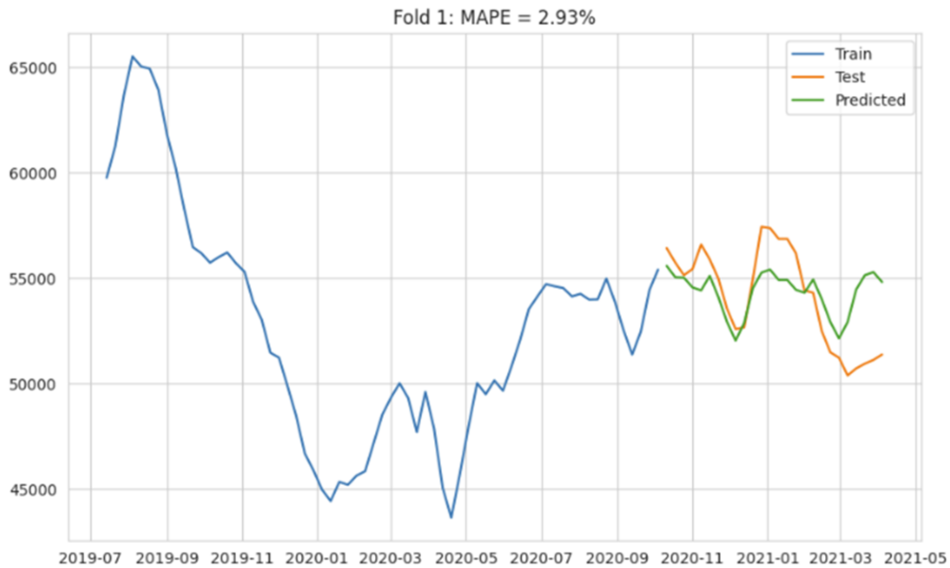

Note. Initial cross-validation fold used to evaluate the accuracy of the SARIMA(1,0,1)(1,1,1)12 model for total pallet positions across all customers. The actual data is represented by the orange line, while the forecast is depicted by the green line for total pallet positions across all customers.

The results presented in Table 2 display a similar view of the top 5 AIC scores for our models, but this time sorted based on “Mean MAPE”. This highlights that although the AIC score is a useful measure for selecting the best model, it is not always indicative of the best performance. In fact, as shown in the table, the model with the lowest AIC score did not have the lowest mean MAPE. The best model had an Average MAPE of 5.28%.

Table 2

Top 5 SARIMA Models: Full Analysis

Combination	p	d	q	P	D	Q	m	AIC	p_values < 0.05	Mean MAPE	SD MAPE
1	1	0	1	1	1	1	12	2,784	Yes	5.28	2.09
2	1	0	2	1	1	1	12	2,763	Yes	5.39	2.20
3	1	0	0	2	1	1	12	2,624	Yes	5.69	1.73
4	2	0	0	1	1	0	12	2,777	Yes	5.90	2.23
5	2	0	0	2	1	0	12	2,567	Yes	6.01	2.13

Note. Table providing top 5 SARIMA models sorted by best (smallest) mean MAPE score for total pallet positions across all customers. Additionally, the table displays the standard deviation of MAPE scores for these models. It also indicates whether the p-values for the models were below 0.05 or not.

Once we found a model with all significant coefficients, we conducted a residuals analysis to assess its adequacy. We checked for normality and autocorrelation in the residuals. Figure 24 shows the residual analysis results for the SARIMA (1,0,1)(1,1,1)12 forecast across the entire customer base. The residuals displayed a normal distribution around the mean, indicating that our model captured the underlying data patterns. We observed no significant autocorrelation in the residuals. We also confirmed the distribution of residuals by plotting a histogram with a mean close to zero. Despite some models failing to meet these criteria, we retained them for their significance, recognizing their potential limitations in fully representing the dataset's behavior.

Figure 24

Residual Analysis: SARIMA(1,0,1)(1,1,1)12

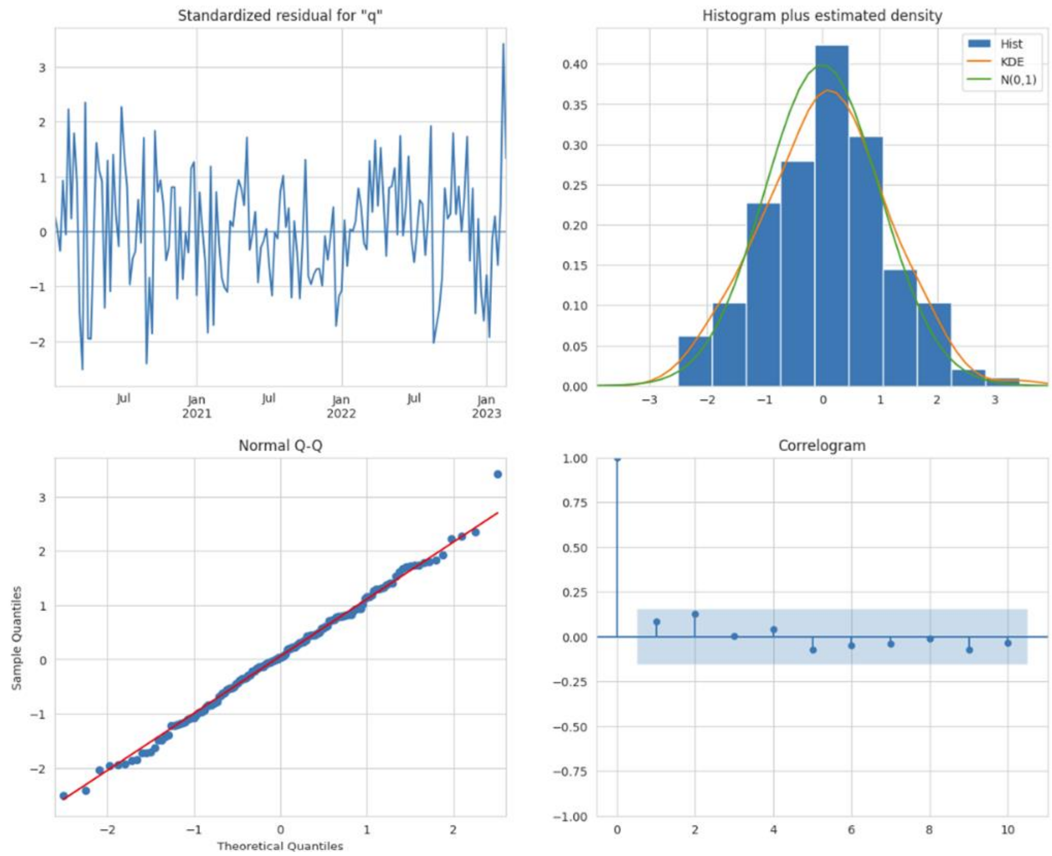

Note. Plots show overall residual analysis (PACF, Normal Distribution and Scatter Plot) for total pallet positions across all customers.

We performed the above analysis on all data aggregation groups discussed above resulting in 3,024 SARIMA models created. We chose the best SARIMA with m=12 and SARIMA with m=52 in our analysis for choosing the best model overall. This table can be found in Table A in the Appendix below.

3.9. Forecasting with Prophet

To utilize Facebook Prophet for forecasting, a data frame with two columns is required. The first column, labeled "ds," contains the date or datetime column of the time series data and must be in a specific format, such as YYYY-MM-DD for daily data or YYYY-MM for monthly data. The second column, named "y," contains the target variable or response variable of the time series data that we want to forecast.

After several iterations of parameter adjustments, we selected four models with the highest MAPEs across different data segmentations. The first model chosen was the baseline Prophet model, which has two default settings that are relevant to this capstone. The first default setting assumes that the data has a linear trend, which is appropriate for this capstone. The second default setting allows for up to 25 changepoints in the linear trend. Prophet uses a Bayesian changepoint detection algorithm to detect changes in linear trends, meaning that the linear trend can change up to 25 times across the dataset. Figure 25 displays the actual versus predicted values (y_hat) for the total site inventory at the weekly level, and also illustrates the trend and seasonality components of the model. In this model, the trend changed four times.

Figure 25

Prophet Forecast: Baseline

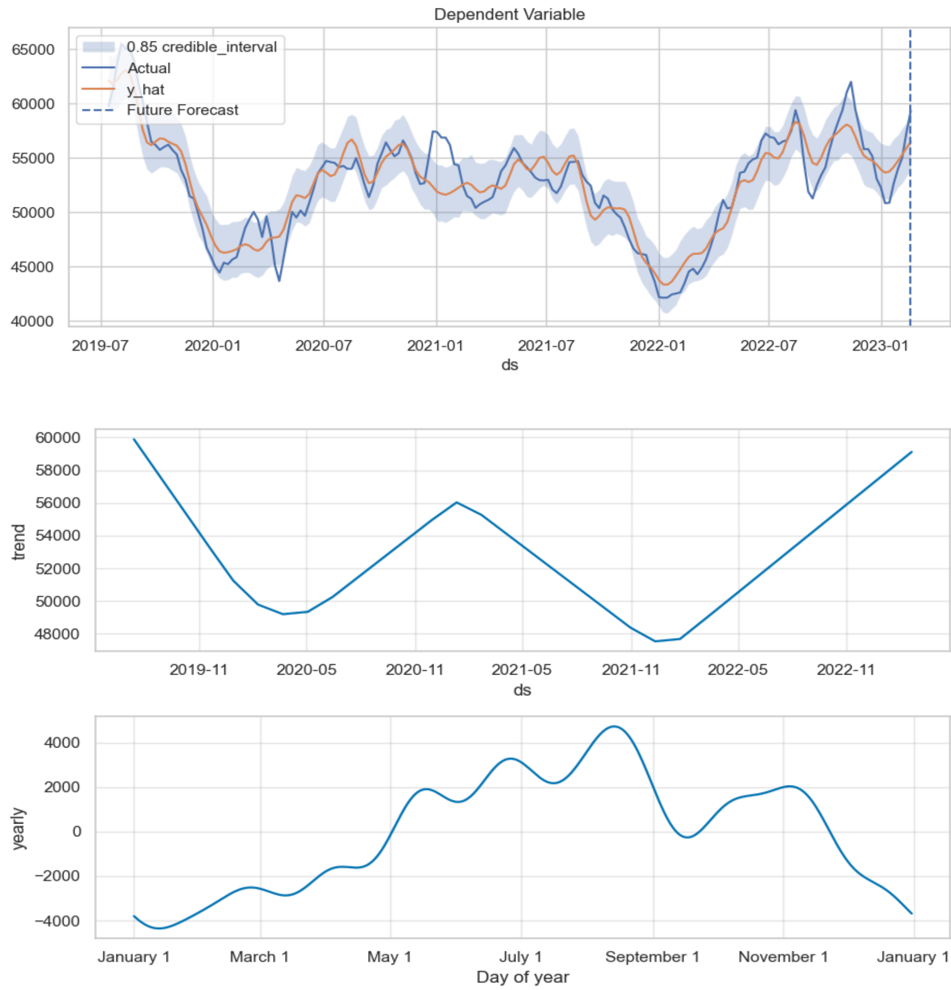

Note. Plots illustrate the model fit, trend, and yearly seasonality of the baseline model of Prophet for total pallet positions across all customers.

The next model we found useful only allowed one changepoint, as having too many changepoints can make the model erratic. By reducing the number of changepoints, we achieved better results on certain data segments. Figure 27 shows the comparison of the actual values with the predicted values (y_hat) for the total site inventory at the weekly level, and also displays the trend and seasonality components of the model. In this case, the trend changes only once throughout the entire dataset.

Figure 26

Prophet Forecast: Changepoint = 1

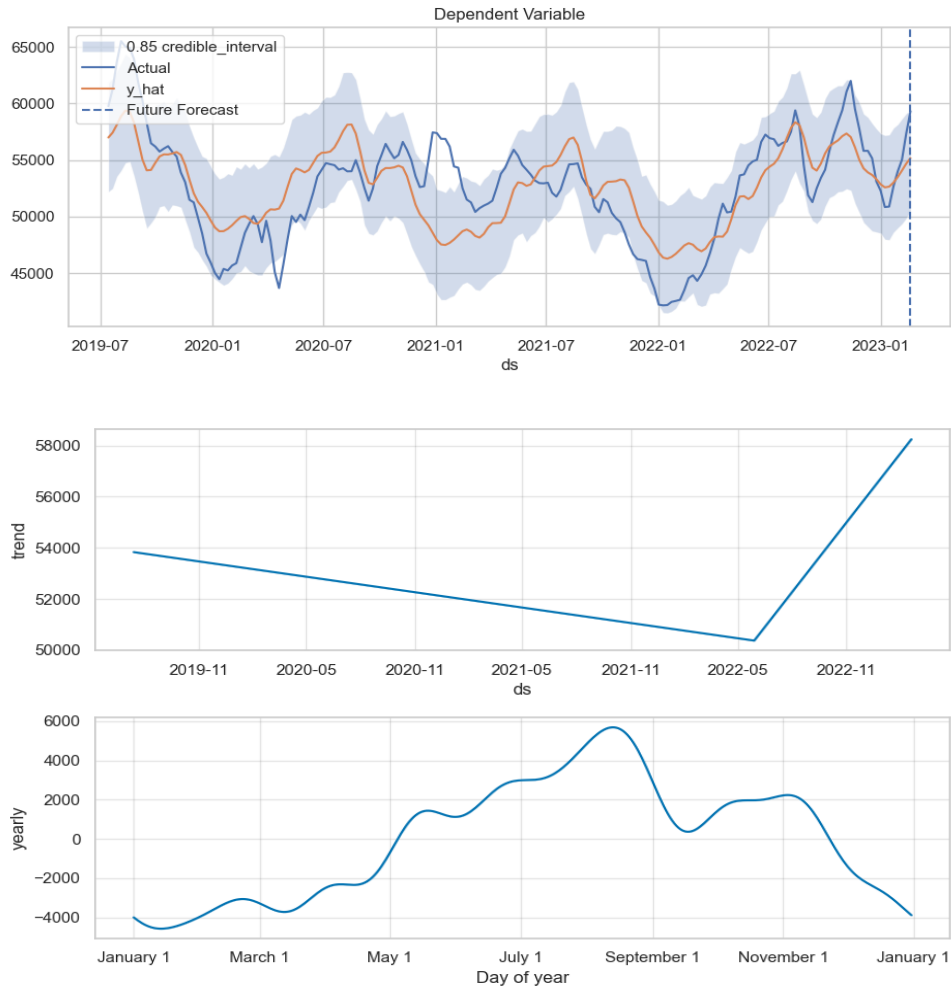

Note. Plots illustrate the model fit, trend, and yearly seasonality of the Prophet model with the "n_changepoint" parameter equal to 1 for total pallet positions across all customers.

The next model we found useful was to not allow any changepoints meaning that the prophet model fitted one linear line across the entire dataset. Figure 27 shows the comparison of the actual values with the predicted values (y_hat) for the total site inventory at the weekly level, and also displays the trend and seasonality components of the model. In this case, the trend does not change throughout the entire dataset.

Figure 27

Prophet Forecast: Changepoint = 0

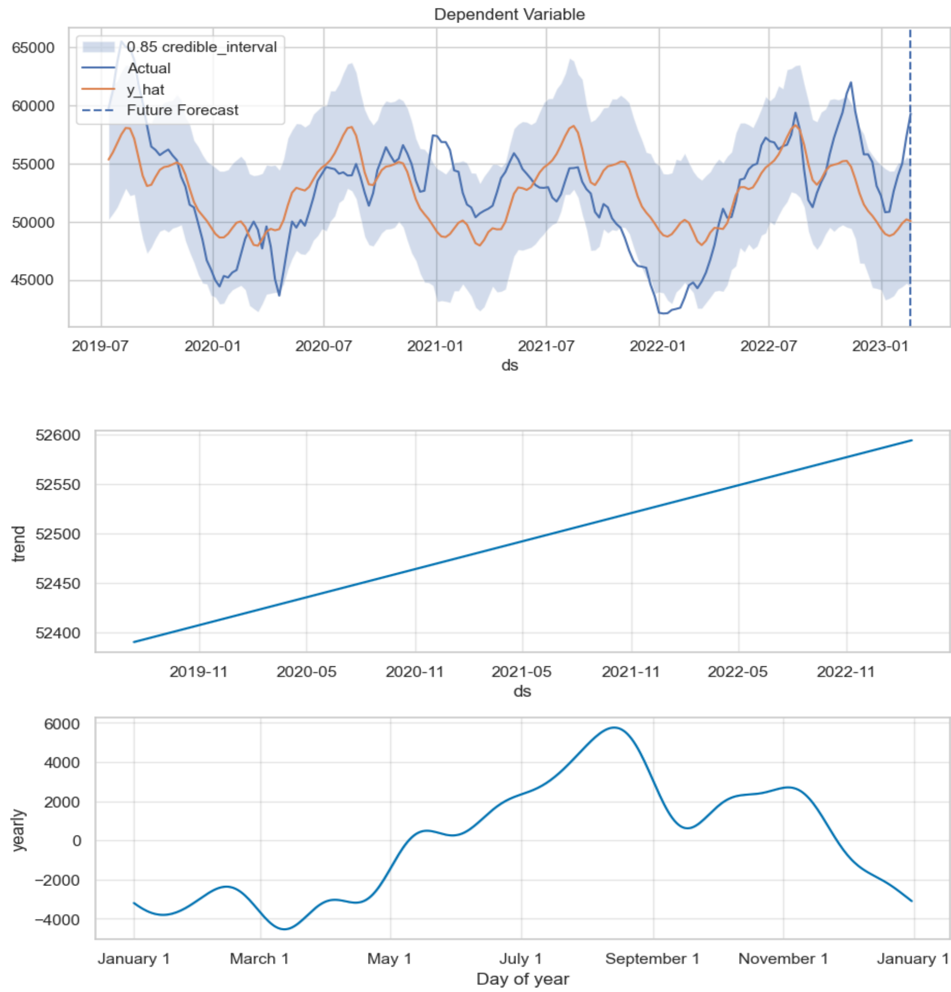

Note. Plots illustrate the model fit, trend, and yearly seasonality of the Prophet model with the "n_changepoint" parameter equal to 0 for total pallet positions across all customers.

The final model we found to be useful involved no trend, indicating that there was no growth over time. We observed that this model was the best fit for several data segments when using Prophet. We will further elaborate on this finding in our results and discussion section. Figure 28 presents a comparison between the actual and predicted values (y_hat) for the total site inventory at the weekly level, along with the trend and seasonality components of the model. Notably, there is no discernible trend throughout the entire dataset.

Figure 28

Prophet Forecast: Growth = Flat

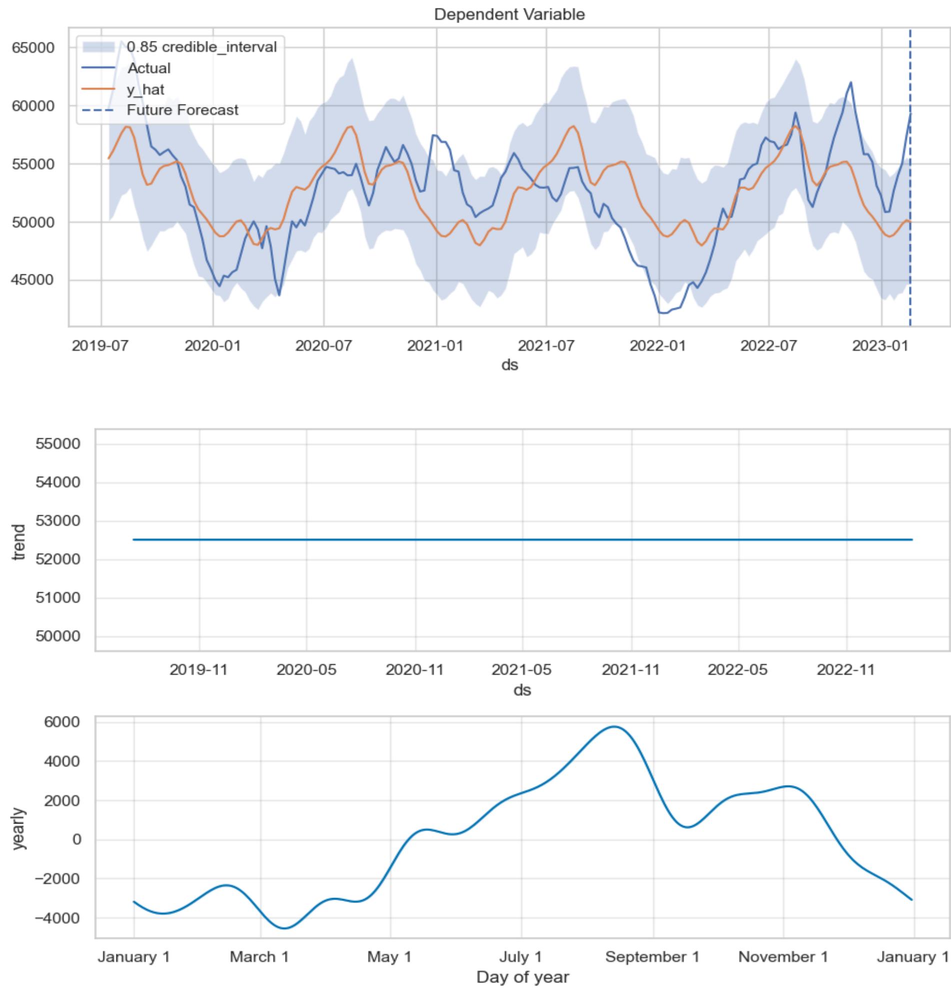

Note. Plots illustrate the model fit, trend, and yearly seasonality of the Prophet model with the "growth" parameter equal to "flat" for total pallet positions across all customers.

We assessed the performance of each of the four Prophet models described above through cross-validation and residual tests, similar to the SARIMA models discussed in section 3.8. All of the Prophet models were included in our analysis to determine the best overall forecasting model, and the results are presented in Table A in Appendix 1.

4. RESULTS

In the methodology section, we provided a comprehensive overview of all the steps involved in the forecasting process. In the Results and Discussion section, we present the most effective forecasting methods for each aggregation performed. We begin by analyzing the most aggregated data, which includes all customers aggregated at the weekly level. This is followed by forecasts for all customers in the freezer and cooler categories. Lastly, we examine forecasts by customer segments created. As we go through the results, we observe that the higher the level of aggregation, the better the forecast accuracy, but this comes at the expense of losing visibility into the individual components that make up the overall forecast. Therefore, it is important to strike a balance between the level of aggregation and the level of granularity in the forecasts.

We introduce two additional features with the first being a confidence interval. The confidence interval chosen for the forecast models is 85%. An 85% confidence interval is a statistical measure that indicates that there is an 85% likelihood that the true value of a population parameter falls within the specified interval. In other words, if the same population is sampled repeatedly and a confidence interval is constructed for each sample, then 85% of the intervals will contain the true population parameter value, while 15% of the intervals will not.

We also introduce inventory capacity constraints into the model. The graphs display two types of constraints. The first is maximum inventory capacity and the second is target utilization. It is important for a site to avoid reaching total capacity as it can lead to issues such as productivity. Hence, based on factors such as depth of racking, sites aim for an inventory level that allows for flexibility. This ideal inventory capacity represents the target capacity for the site.

4.1. Forecast: All Customers

Table 3 shows that SARIMA m(12) (1,0,1) (1,1,1) produced the most accurate forecast for the total site inventory aggregated weekly. This model incorporates a 12-week seasonal pattern and various autoregressive and moving average orders. It assumes that the non-seasonal component of the time series is stationary (d=0), and the seasonal component requires one differencing to become stationary (D=1). The forecast depicted in figure 29, had a mean MAPE of 5.28% and a standard deviation of 2.09% as determined by cross-validation. The 26-week forecast, which covers the period from February 2023 to August 2023, has a flat trend while oscillating around a 12-week seasonal pattern. The upward confidence interval reaches target capacity in week 23 (6-August-2023), which indicates that there is a 7.5% chance that the site reaches capacity during this week.

Table 3

All Customers: Top 3 Models

Forecast Type	Mean MAPE	MAPE Std
SARIMA m(12) (1,0,1) (1,1,1)	5.28%	2.09%
Prophet: No Growth	5.87%	3.11%
Prophet: Baseline	6.05%	3.50%

Note. Top 3 models based on lowest Mean MAPE score for total pallet positions across all customers

Figure 29

All Customers Forecast: SARIMA m(12) (1,0,1) (1,1,1)

Note. Time series showing a 26-week forecast for total pallet positions across all customers. The orange line provides the actual forecast, and the grey shaded area provides the 85% confidence interval of the forecast.

4.2. Forecast: Cooler Customers

As shown in Table 4, SARIMA m(52) (0,1,0) (1,0,0) created the best predictive model for all cooler customers aggregated at the weekly level. The model includes a yearly (52 weeks) seasonal pattern with a seasonal autoregressive order of 1. Additionally, it assumes that the seasonal component of the time series is already stationary (d=0), and the non-seasonal component requires one differencing to become stationary (D=1). The forecast model depicted in figure 30 had a mean MAPE of 11.73% and a standard deviation of 4.21% as determined by cross-validation. The 26-week forecast indicates that the inventory level will be linear with a low negative slope. The figure also illustrates a consistent pattern in both actual and forecasted data, revealing a persistent underutilization of cooler pallet positions. Lastly, the capacity in the figure shows a reduction in late 2022. In late 2022 one of the rooms within the site was changed from cooler to freezer to provide more pallet capacity for freezer customers.

Table 4

Cooler Customers: Top 3 Models

Forecast Type	MAPE Mean	MAPE Std
SARIMA m(52) (0,1,0)(1,0,0)	11.73%	4.21%
SARIMA m(12) (1,0,0)(1,1,0)	11.92%	5.92%
Prophet: One linear Changepoint	15.33%	8.62%

Note. Top 3 models based on lowest Mean MAPE score for total pallet positions across all cooler customers.

Figure 30

Cooler Customers: SARIMA m(52) (0,1,0)(1,0,0)

Note. Time series showing a 26-week forecast for total pallet positions across all cooler customers. The orange line provides the actual forecast, and the grey shaded area provides the 85% confidence interval of the forecast.

4.3. Forecast: Freezer Customers

As shown in Table 5, the Prophet model with the parameter "growth = flat" was the best predictive model for the all freezer customers aggregated at the weekly level. This model showed seasonality from +4,000 to -4,000. The forecast model depicted in figure 31 had a mean MAPE of 6.11% and a standard deviation of 1.87% as determined by cross-validation. The 26 weeks forecast begins with the last fitted point for the model which accounts for the variance between the last actual data point and the first forecast. The overall figure highlights several instances where the target utilization of positions for the freezer was surpassed. Moreover, the forecast begins at a lower baseline and is currently experiencing a seasonal uptrend, with the tail of the forecast confidence interval nearly reaching the target pallet positions. This observation suggests that the freezer customers may require additional inventory capacity in the future.

Table 5

Freezer Customers Top 3 Models

Forecast Type	MAPE Mean	MAPE Std
Prophet: No Growth	6.11%	1.87%
Prophet: Baseline	6.76%	3.74%
Cumulative	7.14%	3.82%

Note. Top 3 models based on lowest Mean MAPE score for total pallet positions across all freezer customers.

Figure 31

Freezer Customers: Prophet Forecast Growth = Flat

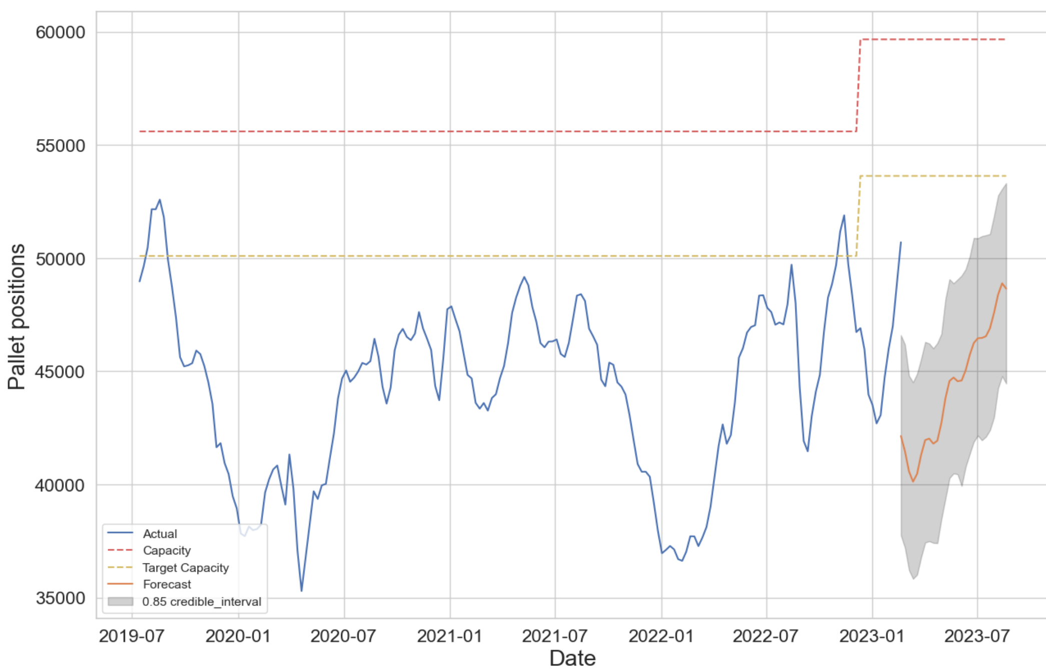

Note. Time series showing a 26-week forecast for total pallet positions across all freezer customers. The orange line provides the actual forecast, and the grey shaded area provides the 85% confidence interval of the forecast.

4.4. Forecast: Customer Segmentation

We first looked at all customer segments associated with cooler business. Table 6 displays the best models obtained for each cooler customer segment. SARIMA models exhibited the best MAPE for each customer segment.

Table 6

Top Models: Cooler Customer Segments

Customer Segmentation	Forecast Type	MAPE Mean	MAPE Std
Low_Inv-High_CV_COOLER	SARIMA m(52) (0,1,0)(1,0,0)	20.62%	4.90%
Low_Inv-Low_CV_COOLER	SARIMA m(12) (1,0,1)(1,1,0)	17.79%	11.17%
MVC Program_COOLER	SARIMA m(12) (1,0,1)(0,1,1)	11.86%	8.93%

Note. Top models for each customer segment on lowest Mean MAPE score for all cooler pallet positions within the site.

Figure 32 provides a graphical representation of actual and forecasted data for each customer segment with the blue dotted line providing the split between the actual data and the forecasted data. It is important to note that these forecasts do not provide the confidence intervals.

Upon analyzing Figure 32, it is further evident that the cooler business has not shown growth over the last few years. The "Low_Inv-Low_CV_Cooler" segment, in particular, has driven the flat to negative growth trend that persists in the forecast. This lack of growth highlights the potential need for strategic interventions to improve the performance of this customer segment or find new customers that could help fill the inventory positions available. In conclusion, Figure 32 provides us with a more detailed understanding of the cooler business, which can be used to develop effective business strategies.

Figure 32

Cooler Customer Forecast by Segment

Note. Time series showing the actual and a 26 week forecast for each customer segment within cooler. The blue dotted line separates the actual data from the forecasted data.

Figure 33 provides the sum of all forecasts within the cooler segments which provides a similar view of the overall forecast viewed in section 4.2 when forecasting of all aggregated cooler customers. This gives more confidence to the idea that the actual pallet positions within the cooler customer base will be flat or declining in the future.

Figure 33

Consolidated Cooler Customer Segment Forecasts

Note. Time series showing the 26-week forecast based on the sum of the top models for each customer segment within cooler. The orange line provides the summed forecasts, and the grey shaded area provides the 85% confidence interval of those forecast.

Next, we analyzed the customer segments related to the freezer business. Table 7 displays the best models obtained for each customer segment within the freezer with SARIMA models exhibiting the best MAPE in all cases except for "MVC Program_FREEZER".

Table 7

Top Models: Freezer Customer Segments

Customer Segmentation	Forecast Type	MAPE Mean	MAPE Std
Key Customer 1	SARIMA m(12) (2,1,1)(0,1,1)	13.17%	6.29%
Key Customer 2	SARIMA m(12) (1,0,2)(0,1,1)	12.84%	5.20%
Key Customer 3	SARIMA s(52) (1,0,1)(0,1,1)	14.55%	3.99%
Low_Inv-High_CV_FREEZER	SARIMA m(52) (1,0,1)(1,1,1)	9.90%	2.29%
Low_Inv-Low_CV_FREEZER	SARIMA m(52) (1,0,1)(1,0,0)	16.86%	6.32%
MVC Program_FREEZER	Prophet: Baseline	7.65%	2.00%

Note. Top models for each customer segment on lowest Mean MAPE score for all freezer pallet positions within the site.

In section 4.3, it was established that the overall growth in the freezer customer segment was flat. Performing the forecast for each customer segment individually uncovered some compelling trends and patterns, as demonstrated in Figure 34.

A significant finding was that the "MVC Program_FREEZER" customer segment is projected to experience steady growth over the next 6 months. On the other hand, the remaining customer segments are mostly expected to remain flat or decline. These forecasted trends suggest that targeted interventions might be needed to stimulate growth in the freezer customer segments that are expected to remain flat. Overall, this analysis reveals important insights that can be leveraged to develop effective business strategies.

Figure 34:

Freezer Customer Forecast by Segment

Note. Time series showing the actual and a 26 week forecast for each customer segment within freezer. The blue dotted line separates the actual data from the forecasted data.

Figure 35 offers a unique perspective by displaying the total of all predictions within the freezer segments, setting it apart from the combined freezer forecast. Instead of showing flat growth similar to section 4.1, this forecast illustrates an upward trend with the "MVC Program_FREEZER" segment driving the majority of the expected inventory growth. Notably, the MAPE for "MVC Program_FREEZER" was only 1.5% greater than the aggregated freezer forecast.

It is important to acknowledge that the combined actual forecasts exceed the targeted capacity, and the confidence intervals surpass the overall capacity. Therefore, if Americold grows inventory levels like the forecasts suggest, it would need to execute one of three strategies. Firstly, Americold could repurpose another room from cooler to freezer if the room sizes align with freezer and cooler forecasts. Secondly, Americold could cannibalize certain customers to free up incremental space for growth in the freezer segment by utilizing 2*2 matrix created in Figure 6. Lastly, Americold could expand the building based on further analysis of future projections. These options would require careful consideration of various factors such as cost, customer impact, and long-term sustainability which should be accounted for in the S&OP process.

Figure 35:

Consolidated Freezer Customer Segment Forecasts

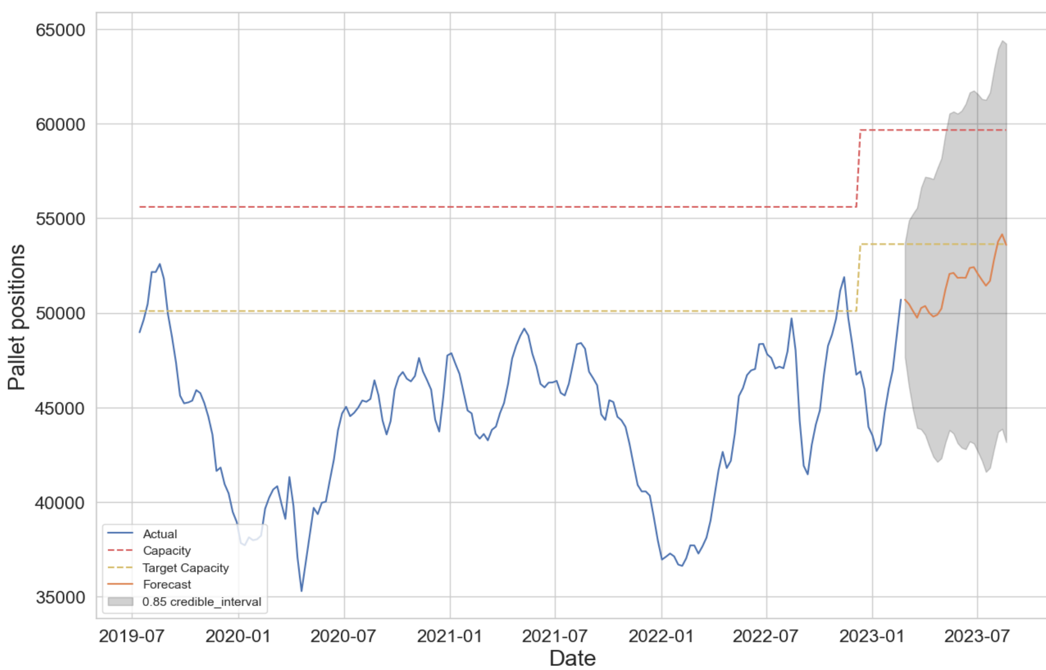

Note. Time series showing the 26-week forecast based on the sum of the top models for each customer segment within freezer. The orange line provides the summed forecasts, and the grey shaded area provides the 85% confidence interval of those forecast.

Lastly, we conducted an analysis of all customer segments to obtain a holistic view of the site. The forecast for the total site inventory aggregated weekly showed flat growth. However, as depicted in figure 36, when we combined all customer segment forecasts, we observed a slight upward trend.

The sum of the actual forecast did not exceed the site's capacity, although the confidence intervals indicated a potential break of the target capacity and overall capacity. However, due to the wide range of these intervals, the likelihood of this occurring within the next 6 months is low.

Figure 36

Consolidated Full Site Customer Segment Forecasts

Note. Time series showing the 26-week forecast based on the sum of the top models for each customer segment across the entire site. The orange line provides the summed forecasts, and the grey shaded area provides the 85% confidence interval of those forecast.

5. DISCUSSION

The role of forecasting in the S&OP process cannot be overstated, as it provides the foundation for a robust plan. As highlighted in the Results section, forecasting enabled us to gain valuable insights into the need to create additional freezer capacity at the site. Furthermore, we identified underutilized space in the cooler that could potentially be repurposed to increase freezer capacity. However, it is worth noting that a complete S&OP process involves more than just forecasting.

In the upcoming discussion section, we will delve into the necessary steps required to develop a robust S&OP process. A crucial aspect of this process is the involvement of a cross-functional team from different departments such as sales, operations, finance, and marketing. This collaboration ensures that all relevant perspectives are considered in the S&OP process, and the resulting plan is aligned with the organization's overall objectives. We will also discuss next steps, opportunities and limitations.

5.1. S&OP Process Framework

The objective of this capstone was to establish a dynamic S&OP process for 3PL companies. Although forecasting is an essential component of any Sales and Operations Planning (S&OP) process, a comprehensive S&OP process requires more than just forecasting. It is essential to note that forecasts are not always entirely accurate, especially with the frequent occurrence of black swan events that have severely impacted supply chains recently. The phrase “Forecasts are always wrong!” has been a common sentiment, and it has become even more prevalent with these unforeseen events. However, even without these events, forecasts inherently lack data points required to make precise projections. Moreover, even if a company possessed every data point today, it would be unrealistic to assume that each data point remains the same in the future as inputs into a future forecast.

Therefore, it is crucial to involve subject matter experts in any comprehensive Sales and Operations Planning (S&OP) process. For this capstone, we propose the structure depicted in Figure 36, which enables subject matter expert feedback in the overall process. The objective of this process is to establish a flow of information between data analysis and subject matter expertise who possess an understanding of the nuances of the business that cannot be gleaned from data alone. A crucial aspect of this process is the involvement of a subject matter experts from different departments such as sales, operations, finance, and marketing. This collaboration ensures that all relevant perspectives are considered in the S&OP process, and the resulting plan is aligned with the organization's overall objectives. In this model, the data team concentrates on enhancing models, while taking into consideration recommendation by subject matter experts. This process should be seamless and flexible to support a robust S&OP process.

Figure 37:

Proposed S&OP Process

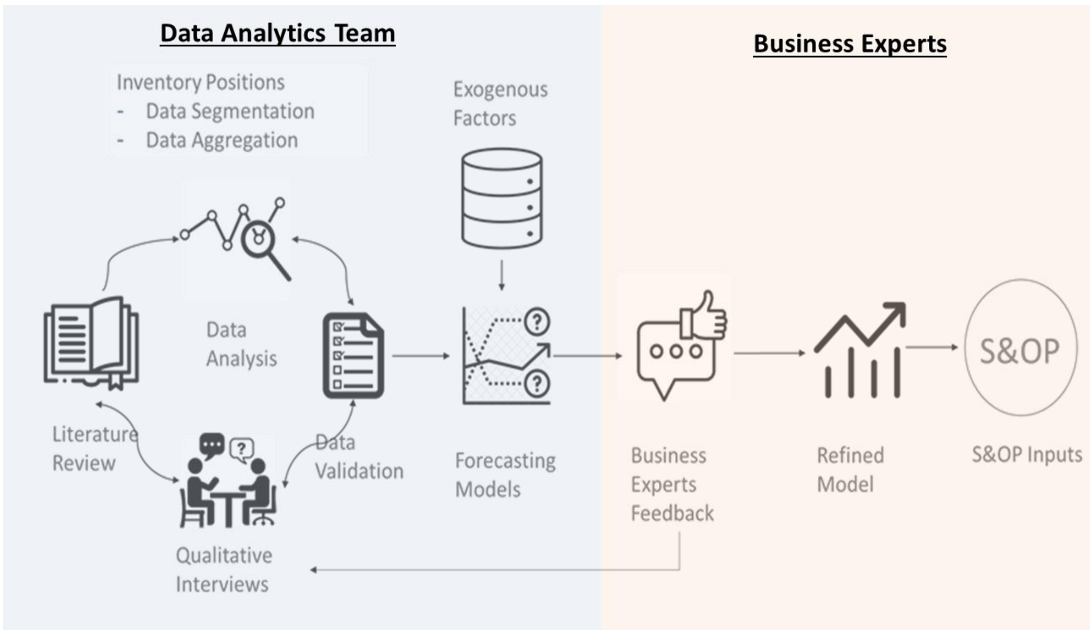

Note. Proposed future process flow for a Robust S&OP process at Americold

It's also important to note that a robust S&OP process contains regular meetings to discuss findings and perform scenario planning. These meetings should allow for the creation and evaluation of different scenarios to account for changes in demand, supply, and other external factors. Discussions pertaining to strategic decisions, such as whether to construct a new site, expand existing infrastructure, or close a site, should be conducted during these meetings.

5.2. Scaling Across Americold's Network

Throughout the Methodology and Results sections, we meticulously designed a framework that facilitates scaling of forecasts across Americold's entire network. The first step involved ensuring that we had access to the precise tables stored in Americold's data lake. To automate the data segmentation, data aggregation, and forecasting steps, we utilized Python. However, running these functions for a single site still required a significant amount of time on local hardware.

To scale this process across Americold's entire network, we recognized the need to run these functions on a local or cloud-based server. By doing so, Americold can significantly reduce the processing time required for each site and efficiently manage the vast amounts of data involved. Additionally, leveraging a server-based solution allows for more advanced data analysis techniques and facilitates collaboration across multiple departments and teams.

Furthermore, we acknowledge the importance of considering the unique challenges and opportunities presented by each site when scaling this framework. While a standardized process can facilitate consistency and efficiency, a tailored approach may be necessary to address site-specific factors and ensure the accuracy and reliability of the forecasts. We suggest Americold roll this out by geographical region to identify any site specific characteristics that do not fall into the framework created.

In conclusion, developing a framework that allows for scaling forecasts across Americold's entire network requires careful consideration of various factors, including data accessibility, processing time, and site-specific challenges. Utilizing a server-based solution and tailoring the approach to each site's needs can enable more accurate and reliable forecasts while optimizing efficiency and scalability.

5.3. **Incorporating Fixed Commitments**

We focused only on actual occupied positions in this model and excluded companies that pay for fixed locations within the site. These fixed locations are purchased to ensure space availability during inventory peaks, which is an industry standard and a key part of Americold's business strategy. Our focus was on understanding the actual pallet positions as forecasting fixed positions would degrade the forecast, though incorporating this information would have enhanced our ability to make informed decisions about the potential adjustments discussed in the Results section. The majority of pallet positions within the site analyzed are currently either occupied or reserved for customers, which does not affect the forecast but has an impact on possible solutions. It's important for the company to track actual pallet positions versus fixed positions as it provides valuable insights for the sales and operations teams. We think it would be beneficial further research to look at binary classification using machine learning to predict the likelihood of a customer decreasing or increasing their fixed commitments in the next contract iteration. This could provide valuable insights to the company and sales team.

5.4. **Self-Regulated Systems**

When analyzing forecasts for the entire site, either fully aggregated by customer or summed over all customer segments, we observed negligible growth. This is because the selected location is a warehouse that operates near its maximum target capacity when considering customers with fixed positions. Managers take measures within the system to prevent inventory from exceeding the target capacity. Such measures may include redirecting goods to another warehouse within the organization or postponing the entry of a new customer. While necessary to ensure great service to customers, this self-regulating process has muted growth trends within the forecast that could have provided valuable insights. Whether coincidence or by choice, the “MVC program_Freezer” has grown over 10K pallet positions while other customers have reduced pallet positions in the freezer while the overall system stays slightly under the target capacity. The significance of subject matter experts in the S&OP process and customer segment-level forecasting becomes evident when forecasting within a self-regulated system, as it allows for the identification of more precise trends. In contrast, there may be other facilities that are growing due to underutilization, and forecasts for those facilities should be treated differently. To address this issue, we propose the creation of a "Site Maturity" classification that examines the average site utilization over a specified period and reflects any changes that occur at the site. This would allow a company to understand any self-regulated biases in the forecast. We recommend further research into this topic.

5.5. **Further Exogenous Variables**

After examining several exogenous variables such as the Consumer Price Index (CPI), total COVID cases, COVID cases exceeding 1 million, and COVID lockdowns, we found minimal to no correlation between these variables and inventory position. Therefore, we decided not to use models like SARIMAX, which involve exogenous variables. However, we believe that there is potential for further research in this area. For example, industry-specific growth could be a factor to consider. The state of the art suggests that temperature-controlled shipments are expected to double by 2032, which could help improve forecasts. Additionally, aggregating across the region or country based on commodity type could reveal growth trends that could inform forecasts for a specific site. Using internal sales data and customer relationship management (CRM) software, such as Salesforce, could also enhance the accuracy of forecasts and provide valuable insights for the S&OP process. We recommend conducting further research in this area to explore these possibilities

5.6. **Using Inbound and Outbound to Forecast Inventory**

The inventory forecasting models developed in this capstone rely on the weekly average of pallet positions in inventory throughout the study period. We recognize that inventory at a given time 't' is a function of inventory at the preceding time 't-1', plus the inbound at time 't', minus the outbound at time 't'. Consequently, inbounds and outbounds can be represented through time series, each of which may exhibit distinct behaviors that are not captured in our current model. For instance, consider a scenario where there are no inbounds or outbounds at a certain time 't', resulting in that both time ‘t-1 and time 't' have the same inventory. Alternatively, if the inbounds and outbounds are both 'X' at time 't', it would also result in the same inventory for time ‘t-1’ and 't'. Although both scenarios lead to the same inventory results, they have different behaviors than can lead to different forecasts for future points in time. Further research on the effectiveness of using inbound and outbound data to forecast inventory positions should be conducted.

5.7. Adding Further Capacities

In section 1.4 of the report, it was highlighted that there are several capacity constraints that can affect the operation of a site. While the focus of this capstone was on inventory positions, it is essential to recognize that inventory positions may not necessarily be the sole factor that can impact a site's performance. Other factors such as labor availability and equipment capacity could also contribute to the site's constraints and determine its overall performance.

Therefore, it is necessary to conduct further research to identify all potential capacity constraints and their interdependencies to create a comprehensive understanding of the site's characteristics. This research should include forecasting demand and quantifying the capacity required to meet that demand. By doing so, it would be possible to identify potential bottlenecks in the site's operations and take proactive measures to address them.

Moreover, by having a complete understanding of the site's capacity constraints, the company can better allocate resources, optimize their operations, and ultimately improve their overall efficiency and profitability.

6. Conclusion

Our research was initiated by a challenging research question posed by our sponsor company, "How can the company balance operational costs and service levels while meeting both present and future demand?" This question is pertinent not only to our sponsor but also to the broader 3PL industry, thus making our findings applicable on a larger scale.

In order to address the research question posed by our sponsor company, we conducted a comprehensive literature review and discovered that the most suitable approach was to develop an S&OP framework. However, we found a significant gap in research regarding logistics and 3PL companies in relation to S&OP processes. This is due to the fact that 3PLs are situated between manufacturers and retailers, and often lack critical data such as manufacturing and point of sale information, which is necessary to generate accurate forecasts for an S&OP process. To address this gap in research, we developed a comprehensive approach that allowed us to make significant contributions towards answering the research question posed by our sponsor company.

Our first contribution was developing a robust segmentation model that serves various purposes within the S&OP process. First, the model provides more dependable forecasts by creating segments that incorporate the coefficient of variance (CV), which guarantees that data with high variability is aggregated. Second, the model aids managers in making decisions by allowing them to identify the next steps for the site. For example, if the forecasts indicate that the site will run out of inventory capacity, a logical next step would be to remove customers with low inventory and high CV (i.e., Dogs in Figure 6) as needed. Additionally, if all the high inventory and low CV (i.e., Cash Cows in Figure 6) are growing and will soon break capacity, the company can consider expansion.

After segmenting the data, we generated reliable forecasts that provided key information for the company. We conducted 256 distinct forecasts on each customer segment and identified several viable options for the analyzed site. Our forecasts revealed the need for additional freezer capacity, as well as underutilized space in the cooler segments that could be repurposed for expansion. Converting a room from cooler to freezer can add significant benefits to the bottom line for the site but expert judgment is required before any actual decision can be made.

Finally, we recommended an S&OP framework that enables the company to scale efficiently across their 240+ facilities globally, including the integration of subject matter experts and the establishment of a feedback loop for forecasts. As there are numerous sites, we suggested using two parameters for filtering the most critical sites for review. Although this may be challenging, all the necessary components are present to allow the company to scale and create a "Dynamic S&OP Process".

REFFERENCES

Americold. 2020, June 26. *The food supply chain.* [Video]. YouTube. https://www.youtube.com/watch?v=FqLxJywe9Uc&t=37s

Arunraj, N. S., Ahrens, D., & Fernandes, M. (2016). Application of SARIMAX model to forecast daily sales in food retail industry. International Journal of Operations Research and Information Systems (IJORIS), 7(2), 1-21.

Chambers, J. C., Mullick, S. K., & Smith, D. D. (1971). How to choose the right forecasting technique. Cambridge, MA, USA: Harvard University, Graduate School of Business Administration.

Future Marketing Insights. n.d. *Frozen food market snapshot*. https://www.futuremarketinsights.com/reports/middle-east-north-africa-frozen-food-market

Hyndman, R.J., & Athanasopoulos, G. (2018) *Forecasting: principles and practice*, 2nd edition, OTexts: Melbourne, Australia. OTexts.com/fpp2. Accessed on <11/27/2022>

Jayaram, J., & Tan, K. C. (2010). Supply chain integration with third-party logistics providers. International Journal of production economics, 125(2), 262-271.

Lieb, R. (1992) The use of third-party logistics services by large American manufacturers. Journal of Business Logistics, 13(2), 29-42.

McKinnon A.C. (2004). Third party logistics in the food supply chain. In Bourlakis M.A, & Weightman P.W.H. (Eds.). Food supply chain management (pp.165-178). Blackwell publishing ltd.

MITx_MicroMasters_SCM_KeyConcepts.pdf. (n.d.)

Oliva F., & Revetria R. (2008). A System Dynamic Model to SupportCold Chain Management in FoodSupply Chain. WSEAS International Conference on Systems.

Pullman M., & Wu Z. (2021). Food Supply Chain Management (2nd ed.). Routledge.

Sashi, Cerchione R., Singh R., Centobelli P., & Shabani A. (2013). Food cold chain management: from a structured literature review to a conceptual framework and research agenda. International journal of logistics management, 29(3), 792-821.

Smith D., & Sparks L. (2004). Temperature controlled supply chains. In Bourlakis M.A, & Weightman P.W.H. (Eds.). Food supply chain management (pp.179-198). Blackwell publishing ltd.

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. https://doi.org/10.1016/j.ejor.2015.11.010

Taylor, S. J., & Letham, B. (2018). Forecasting at scale. The American Statistician, 72(1), 37-45.

Tuomikangas, N., & Kaipia, R. (2014). A coordination framework for sales and operations planning (S&OP): Synthesis from the literature. International journal of production economics, 154, 243-262.

U.S. Department of Agriculture USDA (2021). *Food market and prices.* https://www.ers.usda.gov/topics/food-markets-prices/

Zhang, Y., Xu H., & Zheng Y. (2012). Chinese residents' cold chain logistics demand forecasting based on GM (1,1) model. African Journal of Business Management, 6(14), 5136-5141.

APPENDIX A: Rank of Each Forecasting Model created

Customer Segmentation	Forecast Type	MAPE Mean	MAPE Std	Model Rank
All Customers	SARIMA m(12) (1,0,1) (1,1,1)	5.28%	2.09%	1
All Customers	Prophet: No Growth	5.87%	3.11%	2
All Customers	Prophet: Baseline	6.05%	3.50%	3
All Customers	Cumulative	6.94%	4.79%	4
All Customers	SARIMA m(52) (1,0,1)(1,1,0)	6.99%	1.33%	5
All Customers	Prophet: Linear Trend	8.48%	4.14%	6
All Customers	Prophet: One linear Changepoint	8.84%	4.38%	7
COOLER	SARIMA m(52) (0,1,0)(1,0,0)	11.73%	4.21%	1
COOLER	SARIMA m(12) (1,0,0)(1,1,0)	11.92%	5.92%	2
COOLER	Prophet: One linear Changepoint	15.33%	8.62%	3
COOLER	Prophet: Baseline	21.71%	11.96%	4
COOLER	Prophet: Linear Trend	23.89%	14.23%	5
COOLER	Prophet: No Growth	24.36%	15.51%	6
COOLER	Cumulative	24.45%	15.09%	7
FREEZER	Prophet: No Growth	6.11%	1.87%	1
FREEZER	Prophet: Baseline	6.76%	3.74%	2
FREEZER	Cumulative	7.14%	3.82%	3
FREEZER	Prophet: One linear Changepoint	7.78%	3.75%	4
FREEZER	Prophet: Linear Trend	8.25%	3.97%	5
FREEZER	SARIMA m(52) (0,1,0)(1,0,0)	11.73%	4.21%	6
FREEZER	SARIMA m(12) (1,0,0)(1,1,0)	11.92%	5.92%	7
Key Customer 1	SARIMA m(12) (2,1,1)(0,1,1)	13.17%	6.29%	1
Key Customer 1	SARIMA m(52) (0,1,1)(0,1,1)	13.91%	6.15%	2
Key Customer 1	Prophet: One linear Changepoint	16.32%	7.57%	3
Key Customer 1	Prophet: Baseline	19.21%	4.62%	4
Key Customer 1	Prophet: Linear Trend	24.91%	11.99%	5
Key Customer 1	Cumulative	36.02%	22.07%	6
Key Customer 1	Prophet: No Growth	36.49%	14.00%	7
Key Customer 2	SARIMA m(12) (1,0,2)(0,1,1)	12.84%	5.20%	1
Key Customer 2	SARIMA m(52) (1,0,1)(0,1,0)	15.11%	3.26%	2
Key Customer 2	Cumulative	16.92%	2.91%	3
Key Customer 2	Prophet: No Growth	16.98%	3.30%	4
Key Customer 2	Prophet: Baseline	18.26%	10.33%	5
Key Customer 2	Prophet: One linear Changepoint	18.39%	10.29%	6
Key Customer 2	Prophet: Linear Trend	18.39%	10.31%	7
Key Customer 3	SARIMA s(52) (1,0,1)(0,1,1)	14.55%	3.99%	1
Key Customer 3	SARIMA m(12) (2,0,0)(1,1,0)	14.58%	3.42%	2
Key Customer 3	Prophet: One linear Changepoint	14.61%	11.24%	3
Key Customer 3	Prophet: No Growth	14.97%	6.48%	4
Key Customer 3	Cumulative	15.19%	5.42%	5

Customer Segmentation	Forecast Type	MAPE Mean	MAPE Std	Model Rank
Key Customer 3	Prophet: Baseline	17.11%	11.19%	6
Key Customer 3	Prophet: Linear Trend	20.54%	13.36%	7
Low_Inv-High_CV_COOLER	SARIMA m(52) (0,1,0)(1,0,0)	20.62%	4.90%	1
Low_Inv-High_CV_COOLER	SARIMA m(12) (2,0,2)(1,0,1)	21.69%	5.56%	2
Low_Inv-High_CV_COOLER	Cumulative	25.09%	10.26%	3
Low_Inv-High_CV_COOLER	Prophet: No Growth	26.50%	10.74%	4
Low_Inv-High_CV_COOLER	Prophet: Baseline	29.18%	10.38%	5
Low_Inv-High_CV_COOLER	Prophet: Linear Trend	31.04%	14.22%	6
Low_Inv-High_CV_COOLER	Prophet: One linear Changepoint	31.04%	14.22%	7
Low_Inv-High_CV_FREEZER	SARIMA m(52) (1,0,1)(1,1,1)	9.90%	2.29%	1
Low_Inv-High_CV_FREEZER	SARIMA m(12) (2,0,0)(0,1,1)	12.38%	3.65%	2
Low_Inv-High_CV_FREEZER	Prophet: Baseline	13.95%	5.53%	3
Low_Inv-High_CV_FREEZER	Prophet: No Growth	17.61%	10.56%	4
Low_Inv-High_CV_FREEZER	Cumulative	19.38%	9.10%	5
Low_Inv-High_CV_FREEZER	Prophet: One linear Changepoint	28.91%	11.94%	6
Low_Inv-High_CV_FREEZER	Prophet: Linear Trend	29.23%	11.33%	7
Low_Inv-Low_CV_COOLER	SARIMA m(12) (1,0,1)(1,1,0)	17.79%	11.17%	1
Low_Inv-Low_CV_COOLER	SARIMA m(52) (0,1,0)(1,0,0)	17.80%	7.87%	2
Low_Inv-Low_CV_COOLER	Prophet: One linear Changepoint	24.40%	15.34%	3
Low_Inv-Low_CV_COOLER	Prophet: Baseline	32.74%	20.77%	4
Low_Inv-Low_CV_COOLER	Prophet: Linear Trend	37.22%	19.06%	5
Low_Inv-Low_CV_COOLER	Cumulative	45.11%	37.62%	6
Low_Inv-Low_CV_COOLER	Prophet: No Growth	45.38%	37.48%	7
Low_Inv-Low_CV_FREEZER	SARIMA m(52) (1,0,1)(1,0,0)	16.86%	6.32%	1
Low_Inv-Low_CV_FREEZER	SARIMA m(12) (1,0,0)(0,1,1)	19.66%	2.38%	2
Low_Inv-Low_CV_FREEZER	Prophet: One linear Changepoint	20.03%	7.38%	3
Low_Inv-Low_CV_FREEZER	Prophet: Linear Trend	21.46%	8.93%	4
Low_Inv-Low_CV_FREEZER	Prophet: No Growth	28.55%	23.32%	5
Low_Inv-Low_CV_FREEZER	Prophet: Baseline	31.66%	14.76%	6
Low_Inv-Low_CV_FREEZER	Cumulative	32.88%	28.47%	7
MVC Program_COOLER	SARIMA m(12) (1,0,1)(0,1,1)	11.86%	8.93%	1
MVC Program_COOLER	Prophet: Linear Trend	14.13%	10.57%	2
MVC Program_COOLER	Prophet: Baseline	16.73%	10.54%	3
MVC Program_COOLER	SARIMA m(52) (0,0,1)(1,0,0)	16.79%	3.94%	4
MVC Program_COOLER	Prophet: One linear Changepoint	18.90%	10.78%	5
MVC Program_COOLER	Cumulative	29.85%	7.93%	6
MVC Program_COOLER	Prophet: No Growth	31.81%	9.83%	7
MVC Program_FREEZER	Prophet: Baseline	7.65%	2.00%	1
MVC Program_FREEZER	Prophet: One linear Changepoint	9.01%	5.15%	2
MVC Program_FREEZER	SARIMA m(52) (0,1,0)(1,0,0)	11.73%	4.21%	3
MVC Program_FREEZER	SARIMA m(12) (1,0,0)(1,1,0)	11.92%	5.92%	4
MVC Program_FREEZER	Prophet: Linear Trend	12.95%	6.60%	5
MVC Program_FREEZER	Cumulative	14.46%	7.54%	6

Customer Segmentation	Forecast Type	MAPE Mean	MAPE Std	Model Rank
MVC Program_FREEZER	Prophet: No Growth	15.17%	6.34%	7

APPENDIX B: Individual Segment Forecasts for Highest Ranked Model

Key Customer 1 SARIMA m(12) (2,1,1)(0,1,1)

Key Customer 2 SARIMA m(12) (1,0,2)(0,1,1)

Key Customer 3 SARIMA m(52) (1,0,1)(0,1,1)

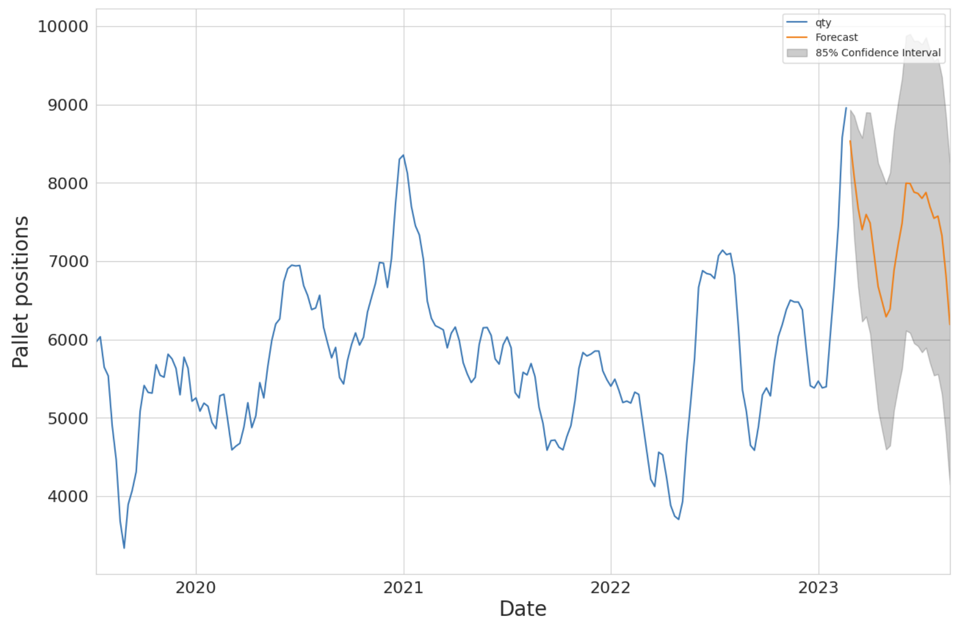

Low_Inv-High_CV_Cooler SARIMA m(52) (0,1,0)(1,0,0)

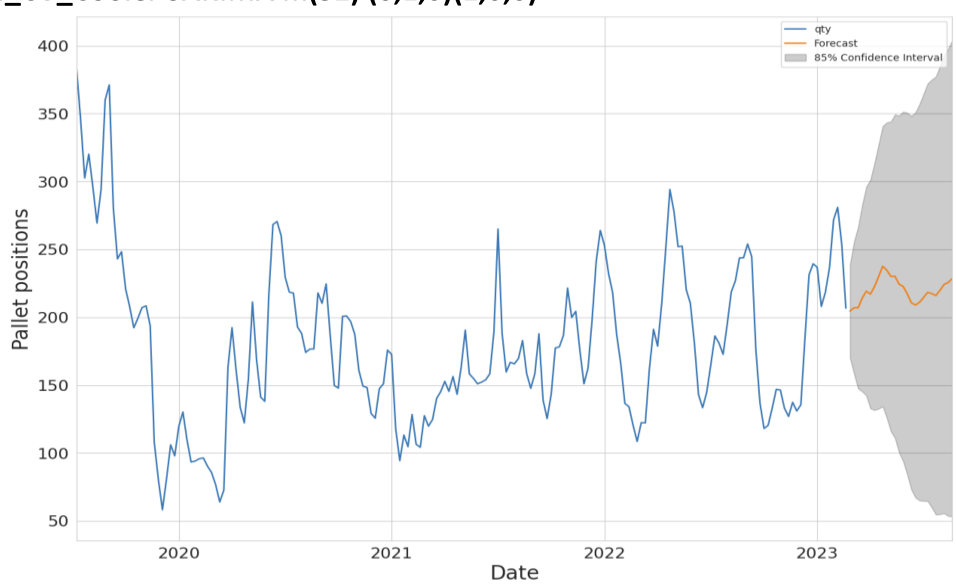

Low_Inv-High_CV_Freezer SARIMA m(52) (1,0,1)(1,1,1)

Low_Inv-Low_CV_Cooler SARIMA m(12) (1,0,1)(1,1,0)

Low_Inv-Low_CV_Freezer SARIMA m(52) (1,0,1)(1,0,0)

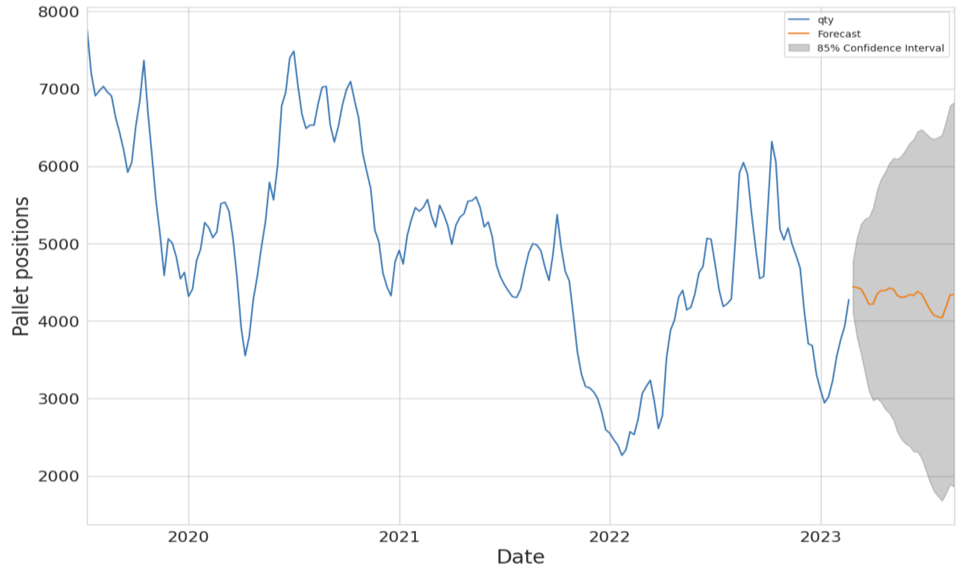

MVC Cooler SARIMA m(12) (1,0,1)(0,1,1)