Designing Inventory Management Strategy for a Fill Rate of 98%

by

Ashish P Chhabria

Bachelor of Commerce (B.Com.), H.R. College of Commerce & Economics, India

and

Deviana Sia

Bachelor of Business Administration (BBA), Wilfrid Laurier University, Canada

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

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

May 2022

Signature of Author:____
Department of Supply Chain Management
May 6, 2022

Signature of Author: ____
Department of Supply Chain Management
May 6, 2022

Certified by: ____
Dr. Josué C. Velázquez Martínez
Director, Sustainable Supply Chain Management & LIFT Research Labs
Capstone Advisor

Certified by: ____
Dr. Ilya Jackson
Postdoctoral Associate, Center for Transportation and Logistics
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

Designing Inventory Management Strategy for a Fill Rate of 98%

by

Ashish P Chhabria

and

Deviana Sia

Submitted to the Program in Supply Chain Management
on May 6, 2022 in Partial Fulfillment of the
Requirements for the Degree of Master of Applied Science in Supply Chain Management

ABSTRACT

Order fill rate is a critical performance metric in retail supply chain operations. Retailers use it to ensure deliveries received from their suppliers are in full quantities as per order. The retailers levy fines on suppliers that fail to comply with the metric. C.H. Robinson has a division that provides retail consolidation services to multiple suppliers. It arranges consigned inventory from multiple suppliers, stores it, and ships it to retailers in full truckloads as per order demand. They are interested in designing an inventory strategy that ensures a 98% order fill rate, thereby minimizing fines charged by retailers. An inventory strategy is focused on three key aspects i.e., optimal review interval, order quantity, and safety stock requirement. This project uses historical order and inventory data provided by C.H. Robinson to design an inventory strategy. The methodology taken is to narrow the focus down to 50 top-selling SKUs out of a total of 3,769 that consistently represent a significant share of the total shipments out of the distribution center. Upon identification of top-selling SKUs, two steps are taken to build a strong foundation before creating an inventory strategy. A forecast is built using techniques such as autoregressive integrated moving average (ARIMA) and error trend and seasonality (ETS) to ascertain the historical volatility in demand. After which the research uses the forecast accuracy to build optimal inventory levels required to achieve order fill rate targets. Furthermore, SKUs that show similar characteristics in terms of fill rate, volatility, and forecast accuracy are segmented into three clusters using k-means clustering. Thereafter, a periodic review inventory control system is used to obtain the optimal review intervals, order quantity, and safety stock levels for each of the three clusters. The research paper suggests an optimal amount of inventory that C.H. Robinson should hold in its DC to ensure an order fill rate of 98%. It also compares it with existing inventory levels maintained at the DC for each cluster, and the corresponding fill rate performance for each cluster. Ultimately, the research paper explores the trade-off of higher inventory holding costs associated with maintaining inventory levels geared towards achieving a 98% order fill rate performance. The research paper also provides C.H. Robinson with a framework they can use to make the best financial decision, given the trade-off mentioned above.

Capstone Advisor: Dr. Josué Velázquez Martínez
Title: Director, Sustainable Supply Chain & LIFT Research Labs

Capstone Co-Advisor: Dr. Ilya Jackson
Title: Postdoctoral Associate, Center for Transportation & Logistics

ACKNOWLEDGEMENTS

A heartfelt thanks to our amazing capstone advisors **Dr. Josué Velázquez Martínez** and **Dr. Ilya Jackson.** We are grateful for your support, expertise, and motivation throughout our capstone journey.

We'd like to thank **Toby Gooley**, for her valuable and constructive feedback in editing our capstone. We are grateful for your patience and consideration throughout.

To the incredible team at our sponsor company, we appreciate your valuable insights that have helped us design our inventory strategy.

Finally, we would like to thank our classmates, the Center for Transportation and Logistics, and the entire MIT community. We will cherish our comradery for a long time to come.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF EQUATIONS

1. INTRODUCTION

The introduction chapter of the research describes the challenges in managing inventory in the retail sector and barriers to achieving a high order fill rate. After exploring the problems prevalent in achieving fill rate targets in retail, the chapter delves into understanding the business processes and challenges specific to C.H. Robinson.

1.1 Motivation

The consumer-packaged goods (CPG) industry is one of the largest industries in North America. Generally, high variability is observed in CPG's supply chain downstream, mostly due to complex demand patterns, dictated by customer buying behavior and seasonality trends in the market. This leads to ripples in the upstream supply chain resulting in longer lead times, stockouts, and poor responsiveness (Gogineni, n.d.).

Order fill rate is a renowned supply chain service level metric used in retail supply chain operations. It is the mechanism through which retailers ensure the order placed by them are delivered in full, thereby ensuring an optimal supply of inventory. The absence of an effective inventory management strategy can, on the one hand, lead to the inability to meet order fill rate targets, and on the other hand, it can lead to excess inventory. Both excess inventory and failure to meet fill rate targets can lead to a monetary loss.

$1.1 trillion in cash or equivalent i.e., approximately 1.5% of the nominal world GDP is tied up in inventory. Moreover, companies are losing $634.1 Billion each year due to out-of-stocks fines and $471.9 Billion due to excessive stocks, which accounts for 4.1% and 3.2% of total annual revenue for an average retailer(*New Research Report*, 2015).

The ongoing Covid-19 pandemic has further contributed to the complexities of the retail supply chains. Wild fluctuations in demand have destabilized each leg of the supply chain. Inventory levels fell precipitously as retailers delayed purchases during the early months of the crisis. Thereafter, the retailers rushed to get a product from the manufacturer to the customer, but supply chains remained snarled. Furthermore, from ocean freight through the middle and last miles, carriers are experiencing unprecedented congestion that has caused service disruptions (*Best Practices to Optimize the Retail Supply Chain | McKinsey*, n.d.).

To help address the combined challenges of fulfillment cost, service requirements, and productivity improvement, retailers have sought to keep inventories closer to consumption centers. In some cases, this practice has led to higher total inventory in the network. For example, the inventory-turnover ratio at most US department stores has decreased over the past five years as illustrated in **Figure 1** (*The Supply-Chain Solution to Aid Retailers' COVID Recovery | McKinsey*, n.d.). The sluggish inventory turnover and high inventory levels can result in increased capital commitment and holding costs for participants in the retail supply chain.

Figure 1

Increase in Inventory Levels and Reduction in Inventory Turnover in Retail Sector

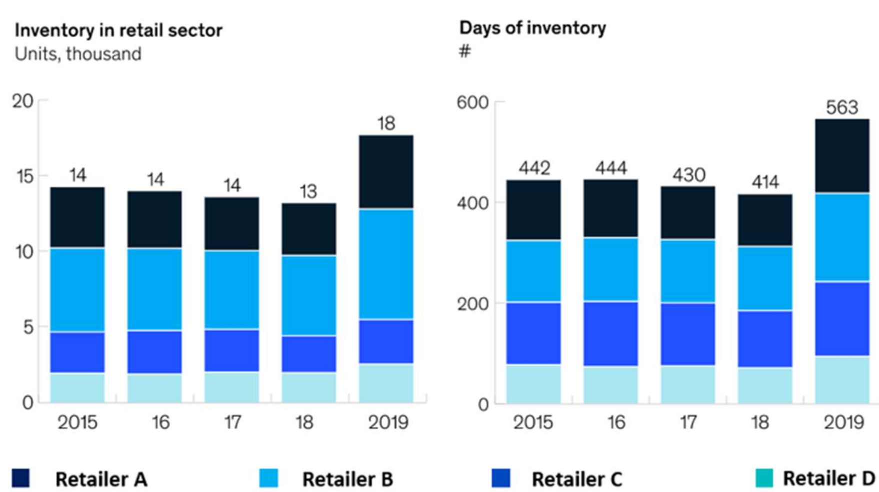

Adapted from McKinsey (*The Supply-Chain Solution to Aid Retailers' COVID Recovery | McKinsey*, n.d.)

Therefore, the two biggest problems in managing inventory in retail, are fines associated with stockouts and costs arising from carrying excess inventory. The next section of the research provides background into C.H. Robinson's retail consolidation division and the challenges they face in meeting retail order demand in full.

1.2 Company Background and Problem Statement

In retail supply chain operations, big retailers use order fill rate metrics to ensure accuracy on the part of their suppliers. They do so, by charging penalization costs if the order fill rate falls below the benchmark of 98%. The lack of an effective inventory strategy can lead to millions of dollars of penalization fines incurred by the suppliers and their service providers.

C.H. Robinson is a company that provides multimodal transportation and third-party logistics as a service. A company division deals with providing third-party retail consolidation services to suppliers. As illustrated in Figure 2, It predominantly arranges consigned inventory from multiple suppliers. Moreover, it receives and ships orders in one full truckload to big retailers with a geographical presence across the United States of America. Consolidating orders into one full truckload leads to several benefits, both for suppliers and retailers. These benefits include improved operational efficiency, reduced carbon footprint, and improved cost efficiency. While C.H. Robinson delivers to many retailers, the biggest is one of the largest retail companies in the world. Managing this key account comes with many challenges.

Figure 2

Retail Consolidation Business Model of C.H. Robinson

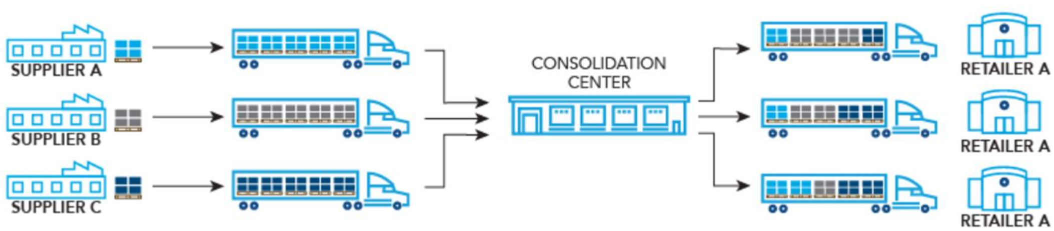

Source: C.H. Robinson website (*Understanding Retail Consolidation | C.H. Robinson*, n.d.)

One of C.H. Robinson's main challenges is a key performance indicator metric (KPI) that retailers implemented, called On Time and In Full (OTIF). OTIF measures the timeliness of orders arriving at retailer distribution centers (DCs) and the fill rate of orders. A big retailer in business with C.H. Robinson expects a 98% OTIF on all orders. The OTIF metric forces supplier organizations to improve operations to avoid fines for failing to meet the 98% compliance threshold (Bower, 2021). These fines, which are in $ millions are then cascaded to C.H. Robinson, which provides consolidation services to these suppliers. This research focuses on designing an inventory strategy, such that the sponsor company can achieve a 98% order fill rate to avoid penalties from falling short on the metric.

At present, C.H. Robinson recommends that suppliers maintain inventory levels worth two weeks of retailer demand at C.H. Robinson's DC. However, this research paper aims to identify a more effective way to manage inventory for the retail consolidator and the suppliers.

Description of C.H. Robinson's Business Processes

A series of interviews were conducted to gain a deeper understanding of the reasons behind the falling fill rate metric and understand the business processes at the C.H. Robinson distribution center. Figure 3 illustrates the process from the receipt of the order to the shipment of products to the retailers.

Figure 3

C.H. Robinson Process Flow

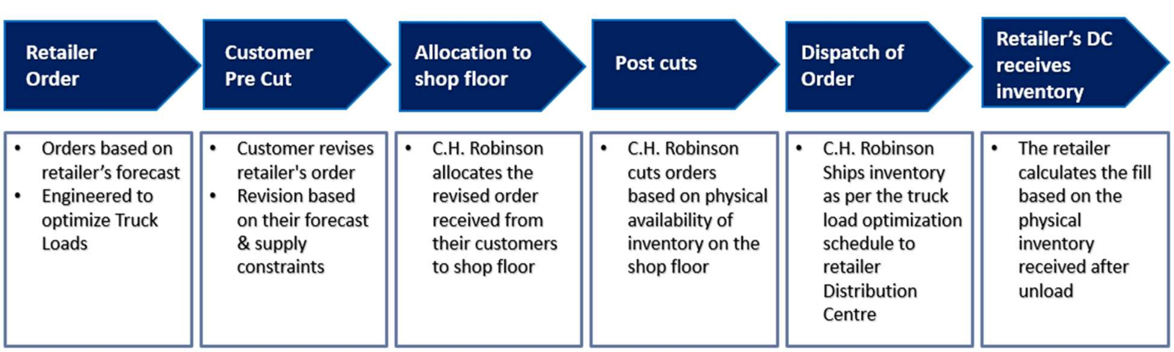

C.H. Robinson's Process with Respect to Inventory Management

As shown in Figure 3, C.H. Robinson receives the order directly from the retailers. Suppliers then revise these orders based on their point of sale forecasts. Once, the order quantity to be shipped is confirmed by the suppliers, C.H. Robinson allocates the order to the shop floor. The orders are picked, packed, and shipped based on the inventory physically available at the distribution center.

Key Supplier's Responsibilities with Respect to Inventory Management

For the research, the biggest supplier of C.H. Robinson was interviewed to understand their role in inventory management

Building a Forecast Based on Point of Sale Demand - The supplier has access to the point of sale demand at the retailers. They base their forecast on that data and revise retailers' orders to that effect. In this way, they make sure they satisfy their end consumer's demand.

Production – Supplier has multiple subcontract manufacturers that ship products directly to C.H. Robinson.

Visibility of Operations – The supplier's supply chain department monitors the inventory on hand at C.H. Robinson, production schedule at manufacturing locations, and forecasting based on the point of sale demand.

List of reasons contributing to the falling fill rate metric

The key takeaways from the interview process established that there was significant variability in demand after the Covid-19 hit. Moreover, certain SKUs were more in demand compared to others. The suppliers were also facing trouble shipping products due to supply constraints instigated by the pandemic.

Impact of Covid-19 Outbreak

As mentioned in section 1.1, supply chain disruptions caused by Covid-19 exacerbated the inherent challenges in the effective management of inventory in the retail sector. The two most prominent factors for the falling fill rate metric due to Covid-19 are mentioned below.

- **Demand Dynamics –**Unexpected surge in demand for products in the wellness category – for example, bath salts.
- **Supply Constraints –** Many raw materials for the supplier's products are imported from other geographies and the international supply chain disruptions have impacted their ability to access raw materials.

C.H. Robinson's business processes and global supply chain disruptions make designing an inventory strategy challenging. To elucidate, as a retail consolidation service provider, the company has no control over the supplier's inventory strategy. Most of the suppliers send inventory per their production schedules without any advance shipping notice. Moreover, the suppliers do not fully exploit C.H. Robinson's IT infrastructure to gauge their inventory levels at the consolidator's DC. In the past few months, the order demand from retailers has increased significantly compared to the historical patterns as well.

To be specific, C.H. Robinson's recent performance on the fill rate metric has been around the low 90s % according to their biggest retailer. This is below the set target of 98% and has resulted in severe fines being charged by this retailer. The monetary fines have necessitated a root cause analysis on the consolidator's part concerning the falling fill rate metric. C.H. Robinson has identified two major root causes for the same. One root cause is the shrinkages in terms of loss and damages while handling inventory in their distribution center. The other root cause is stocking out of products that are high in demand because of supply constraints.

On analyzing the data on a deeper level, the C.H. Robinson team believes that stockouts due to the supply constraints represent the bulk of the problem.

Initially, the research explores the challenges prevalent in managing inventory in the retail sector. Thereafter, a deep dive is taken into the processes and problems involved in retail consolidation at C.H. Robinson. In the next section, the paper aims to explore the research that has already been done with respect to designing an inventory strategy at a retail consolidation center to achieve a fill rate of 98%.

2. LITERATURE REVIEW

Through reviewing relevant literature, we sought to identify the appropriate inventory management strategy for C.H. Robinson. To do so, we first review the concept of retail consolidation and learn more about the business processes of retail consolidators. We then research the role of a robust forecast in building an efficient inventory strategy. Thereafter, we gain a deeper understanding of the appropriate forecasting methods for research purposes. Finally, we review segmentation strategies to identify the SKUs that explain the most about the falling order fill rate metric.

2.1 Retail Third-Party Consolidation

Big retailers run a consolidation program for their suppliers to drive efficiency in their supply chains and minimize their carbon footprint. The way consolidation works is, that many of the retail suppliers consolidate their less-than-truckload (LTL) shipments at a third-party consolidator's distribution center to send full truckload (FTL) shipments to distribution centers of big retailers. Figure 4 shows C.H. Robinson's business model as a retail consolidator.

Figure 4

Retail Consolidation Business Model of C.H. Robinson

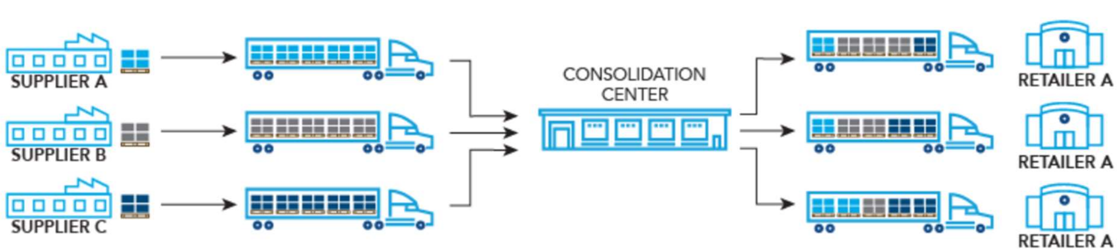

Source: C.H. Robinson website (*Understanding Retail Consolidation | C.H. Robinson*, n.d.)

Retail consolidators such as C.H. Robinson bring multiple benefits to the entire retail supply chain. As the retailers receive full truckload shipments instead of multiple less than truckload shipments, it leads to a reduction in transportation and unloading costs. Moreover, it streamlines material handling, and processing, thereby leading to reduced warehousing costs as well.

2.2 Role of Inventory Management in Retail Consolidation

To ensure efficiency, retailers charge a penalty if the orders are not fully delivered. The way to mitigate the penalty is to ensure appropriate inventory levels at the distribution center. Adequate inventory levels avoid stocking out on products and increase fill rate. It is inventory control systems that enable attaining those adequate inventory levels. They do so, by resolving three issues or problems (Silver et al., 2016):

1. How often the inventory status should be determined?
2. When a replenishment order should be placed?
3. How large the replenishment order should be?

Many inventory control systems have been studied by researchers in the field of supply chain management. However, continuous review and periodic review are the two most often studied. The core difference between the two is the review interval (R). R is the time that elapses between two consecutive moments at which we know the stock level. An extreme case is where there is continuous review; that is, the stock status is always known. Whereas with periodic review, as the name implies, the stock status is determined only every R time unit (Silver et al., 2016). R could be any time unit from a week to a month depending on the nature of the product, demand characteristics, and other business considerations. In practice, a continuous review is very expensive to implement as it entails constant monitoring of inventory levels. Periodic review, on the other hand, allows grouping products with similar demand characteristics and assigning the same review interval to those. It thereby enables workload management on the part of supply chain planners and makes it more reasonable to use. Therefore, in this research, we use the periodic review inventory control system to design an inventory strategy.

2.3 Periodic review, Order-Up-to-Level (R, S) Control Systems

Under a periodic review control system, every R units of time, a replenishment order is placed of enough magnitude to raise the inventory position to the order-up-to-level S. The value of S is determined based on service measure, which for this research is the order fill rate. (Silver et al., 2016)

2.4 Forecasting and its Impact on Setting an Inventory Strategy for C.H. Robinson

Demand forecasting plays a key role in inventory management and determining how profitable a business is (Chawla et al., 2019). Moreover, it is key in determining the order-up-to-level (S) under a periodic review inventory control system. Order-up-to inventory level needs to ensure that the inventory levels can cover the average demand of the product and buffer against any uncertain demand fluctuation in between the review period (R). Therefore, S has two key components, average demand over the review period and the safety stock, also known as buffer stock. The forecast is used to calculate the safety stock component of the order-up-to-level. To do so, the demand for a product is forecasted using the historical demand of the product. Once the forecast is obtained, the root mean squared error (RMSE) of the forecast is used to assess if the forecast explains the variability in demand. (Hyndman & Koehler, 2006) The lower the RMSE, the better it explains the variability. The RMSE is used to calculate the safety stock component of order-up-to-level in the periodic review control system.

2.5 Forecasting Methods

The selection of the forecasting method is a complicated task. It depends on the availability of historical data and the strength of relationships between the forecast variable and any explanatory variables. Furthermore, the way in which the forecasts are to be used (Hyndman & Athanasopoulos, n.d.). Based on the historical data received from C.H. Robinson described in section 3.1, we selected time series forecasting for literature review purposes. Time series forecasting models use mathematical techniques that are based on historical data to forecast demand. It is founded on the hypothesis that the future is an expansion of the past; that's why we can use historical data to forecast future demand. By time series analysis, the forecasting accuracy depends on the characteristics of the time series of demand. If the transition curves show stability and periodicity, we will reach high forecasting accuracy, whereas we can't expect high accuracy if the curves contain highly irregular patterns (Fattah et al., 2018).

Autoregressive integrated moving average (ARIMA) and Error Trend Seasonality (ETS) are traditional statistical models used to model time series. These models are linear since the future values are cramped to be linear functions of past data. Researchers have been focusing much on linear models during the past few decades since they have proved simple in comprehension and application (Fattah et al., 2018).

In the case of ARIMA, it uses observation from historical demand as an input to the equation and uses it to predict future demand. In the case of ETS, the model assumes the time series includes some trend and seasonality in the data. While ARIMA assumes all past data has equal weight, ETS lowers the weight of older data exponentially and puts more weight on the most recent data. This research project has explored both ARIMA and ETS to build a forecast at a stock-keeping unit level (Fattah et al., 2018).

2.6 Evaluating Forecast Accuracy

In time series models, demand forecasts are compared with actual demand, and the difference between the two is known as residual errors. However, the size of the residuals is not a reliable indication of how large forecast errors are likely to be in the future. The accuracy of forecasts can only be determined by considering how well a model performs on new data that were not used when fitting the model(Hyndman & Athanasopoulos, n.d.).

When choosing models, it is common practice to separate the available data into two portions, training and test data, where the training data is used to estimate any parameters of a forecasting method, and the test data is used to evaluate its accuracy. Because the test data is not used in determining the forecasts, it should provide a reliable indication of how well the model is likely to forecast on new data(Hyndman & Athanasopoulos, n.d.). Figure 5 describes the segregation of time series data into training and testing datasets.

Figure 5

Segregation of Time Series Data into Training and Testing Data Sets

Adapted from (Hyndman & Athanasopoulos, n.d.)

A more sophisticated version of training/test sets is time series cross-validation. In this procedure, there are a series of test sets, each consisting of a single observation. The corresponding training set consists only of observations that occurred before the observation that forms the test set. Thus, no future observations can be used in constructing the forecast. Since it is not possible to obtain a reliable forecast based on a small training set, the earliest observations are not considered test sets. Figure 6 illustrates the series of training and test sets, where the blue observations form the training sets, and the red observations form the test sets (Hyndman & Athanasopoulos, n.d.).

Figure 6

Evaluation of Forecast on a Rolling Forecasting Origin

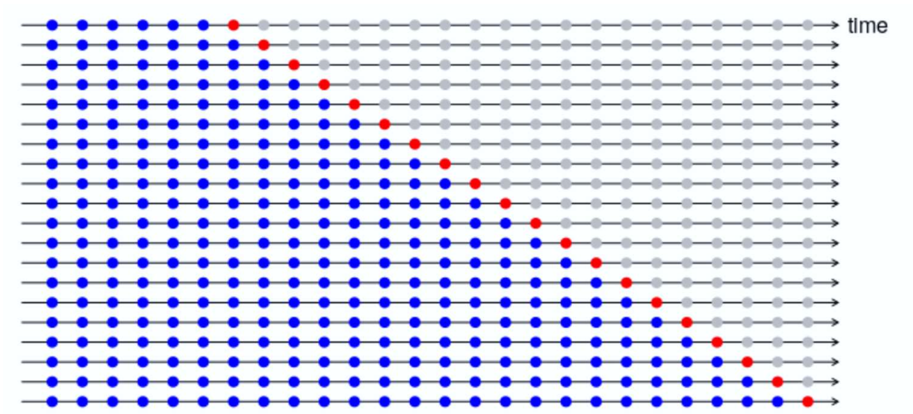

Adapted from (Hyndman & Athanasopoulos, n.d.)

The forecast accuracy is computed by averaging over the test sets. This procedure is sometimes known as "evaluation on a rolling forecasting origin" because the "origin" at which the forecast is based rolls forward in time (Hyndman & Athanasopoulos, n.d.).

Given that the rolling forecasting origin technique is superior in building an accurate forecast, this research project focused on it for cross-validation purposes.

2.7 Selection of Forecasting Technique at a Stock Keeping Unit Level

After reviewing the literature on building a time series forecasting model, and cross-validation, we studied the literature on the selection of forecasting techniques based on forecast accuracy. Forecast errors (e_t) are different than residual errors, mentioned in the section "evaluating forecast accuracy" such that, residuals are calculated on the training set while forecast errors are calculated on the test set. Forecast accuracy can be measured by summarizing forecast errors in different ways (Hyndman & Athanasopoulos, n.d.).

The two prominent ways to measure forecast accuracy that we reviewed are root mean squared error (RMSE) and mean absolute percentage error (MAPE).

Minimizing the RMSE leads to minimizing the average of the forecast errors (e_t), thereby leading to higher accuracy of forecast (Hyndman & Athanasopoulos, n.d.). Furthermore, RMSE has been popular historically because of its theoretical relevance in statistical modeling (Hyndman & Koehler, 2006). In the case of MAPE, it gives a percentage error which has the advantage of being unit-free and therefore is frequently used to compare forecast performances between data sets (Hyndman & Athanasopoulos, n.d.). The MAPE minimizes the percentage of forecast errors (e_t) over the actual observed value of demand (y_t). Therefore, the lowest MAPE would represent the forecast with the highest accuracy.

2.8 Role of Segmentation in Inventory Management

An average inventory system contains an immense number of stock-keeping units (SKUs). In the general case, it is computationally impossible to consider each item individually and manage it under individual inventory policy. Furthermore, the dimensionality of real-world problems requires segmentation of an assortment of SKUs in such a way that, each segment is relatively homogeneous and may be treated under a common inventory policy (Jackson et al., 2019). Therefore, segmentation of inventory is critical to designing an inventory management strategy. In this research paper, we have reviewed a prominent traditional approach to segmentation, "i.e." segmentation based on ABC Analysis. Moreover, we also reviewed a more sophisticated method of segmentation based on an unsupervised machine learning approach, known as k-means clustering.

2.9 Inventory Segmentation based on ABC Analysis

In a perfect world, a company's inventory strategy is not a one-size-fits-all that encompasses all products. On the contrary, it is important to understand the different demand patterns of the products. Thus, a superior inventory strategy comprises sub-strategies for each item such that the strategy fits each item's demand variability characteristics.

In striving for the ideal model, tailoring inventory strategies becomes very challenging as we look at thousands of SKUs. Data processing effort becomes extremely costly as most SKUs vary in behavior. To simplify the problem, we used segmentation. Segmentation is one of the most common methods to aggregate products into groups, specifically, ABC segmentation. Inventory classification using ABC analysis is one of the organization's most widely employed techniques. Normally, the items are classified based on the annual use value, which is the product of annual demand and average unit price. Class A items are relatively few in number but constitute a relatively large amount of annual use value, while class C items are relatively large in number but constitute a relatively small amount of annual use value. (Ramanathan, 2006)

To elucidate, in ABC segmentation, segment "A" is worth 20% of the total SKUs and accounts for 80% of sales. Segment "B" is worth 30% of the total SKUs and accounts for 15% of sales and segment C is worth 50% of the total SKUs and accounts for 5% of the sales (Silver et al., 2017). While this method is useful in determining SKUs importance, it is inefficient for our purpose as it does not take into consideration the demand pattern of the SKUs themselves.

2.10 Inventory Segmentation based on Machine Learning

There are three major categories of machine learning methods: supervised learning, unsupervised learning, and reinforcement learning (Maglogiannis, 2007). Supervised learning methods rely on observations where the correct outcome is known. The algorithm is then trained to get to the correct outcome as close as possible considering variables we think would be relevant to influence the forecast. Unsupervised learning is when the actual result is not given. Finally, reinforcement learning devises a method that rewards desired behavior and/or punishes undesired behavior.

In recent years, more research has been done on a set of promising methods to group stock-keeping units based on cluster analysis. In 2007 the K-means-based SKU segmentation methodology was proposed which is based on cluster analysis. The research aimed to reduce the time required to compute the inventory-control parameters in a large-scale multi-echelon inventory system. Three years later the similar k-means-based approach by Egas and Masel was applied to determine storage assignments. The paper concludes with the statement that the method managed to reduce the number of aisles to retrieve orders by 20–30% compared to a demand-based assignment strategy. Therefore, k-means is a powerful technique based on unsupervised learning to aggregate SKUs (Jackson et al., 2019). Moreover, because the cluster analysis incorporated features with undeniable impact on the inventory management, beyond that utilized by a classical ABC approach, each cluster is homogeneous enough to be treated under a common inventory policy. Thus, the proposed methodology is expected to be efficient for real-world inventory control problems of high dimensionality (Jackson et al., 2019).

Even though our research is not based on the multi-echelon inventory problem, for which it was originally developed, we used k-means clustering. The reason is that k-means is a tool suitable for a large number of SKUs and C.H. Robinson is dealing with thousands of SKUs. Furthermore, it could potentially help group SKUs that represent similar demand and fill rate characteristics, as it clusters similar data sets into groups and determines the underlying pattern (Garbade, 2018).

The k-means algorithm attempts to cluster data separating a set of data observations into k clusters, minimizing the Euclidean distance-based objective function. It repeatedly proceeds in two pivotal steps, namely assignment of each data point to the cluster with the closest centroid and recalculation of the centroids as the mean of all the observations in that cluster until the algorithm converges forming the Voronoi diagram. The number of clusters to be formed using k-means is based on a guess. (Jackson et al., 2019) Given that our research is focused on increasing the order fill rate metric, we selected 3 variables that impact the order fill rate the most “i.e.” variability in demand, forecast accuracy, and order fill rate itself. Since each cluster is homogenous, a common inventory strategy can be devised for SKUs falling under the same cluster.

2.11 Conclusion

While reviewing the literature, we discovered that there has not been much research in the field of inventory management at retail consolidators in specific. Therefore, for our project, we identified literature on aspects of inventory management most relevant to our project and reviewed it.

3. DATA AND METHODOLOGY

Our project aims to develop an inventory strategy for C.H. Robinson to increase the order fill rate metric to 98%. We approached the project in 3 phases as illustrated in Figure 7. In the 1st phase, we identified and gathered the data required to develop an inventory strategy. Thereafter, in the 2nd phase, we conducted our data analysis, built a forecast, segmented SKUs, and developed an inventory strategy. Finally, in the 3rd phase, we validated our inventory strategy with the current framework and made recommendations to the C.H. Robinson team.

Figure 7

Qualitative Process Map of the Capstone Research

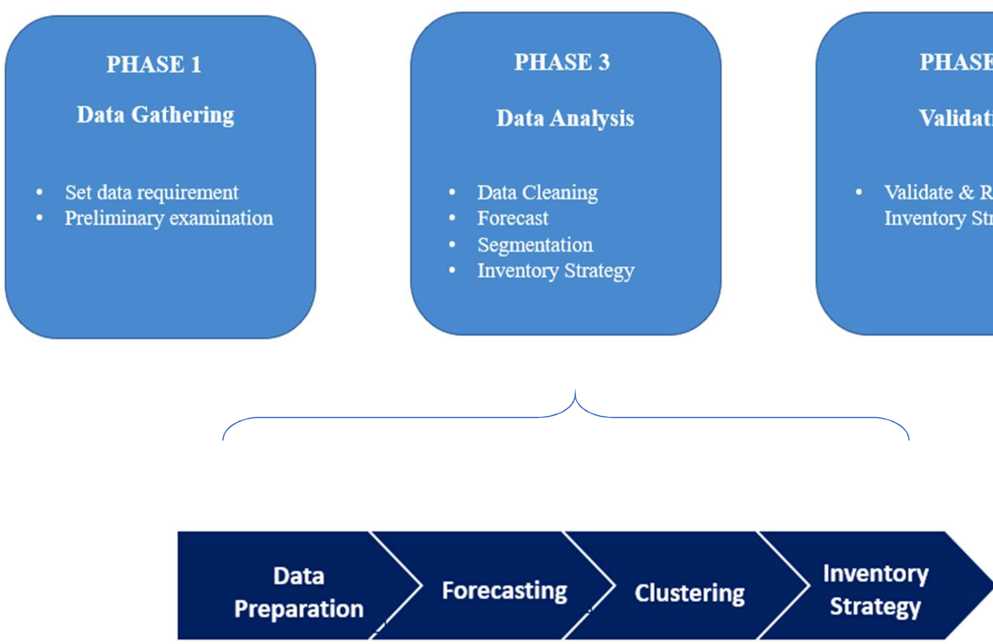

3.1 Data Collection

Based on our understanding of the business processes and the research problem, we identified the relevant data required for our research purposes. Required data was predominantly categorized around the retailer's order data, shipment data, and inventory levels at C.H. Robinson's distribution center. All three categories are fundamental to the purpose of building an inventory strategy.

Time Frame

The time frame selected is for three years. C.H. Robinson, like many other businesses, faced several operational challenges during Covid-19. Therefore, it was considered necessary to get data that represent order patterns for periods before and after the pandemic hit. Therefore, we requested data for the past three years. Furthermore, a time frame of three years provided enough data to train, test, and cross-validate different forecasting methods.

Order Data

We received 11 CSV files, one for each quarter starting from quarter 1 in 2019 up through quarter 3 in 2021. Order files include features of the order number, date, customer number, item number, ship-to location, quantity ordered, quantity shipped, and actual ship date. The order quantity we have chosen for our analysis is the revised order quantity by C.H. Robinson's vendors based on their forecast.

Inventory Data

The inventory data provided the SKU-level inventory on hand per day. For this again we have received 11 CSV files, one for each quarter starting from Quarter 1 in 2019 through Quarter 3 in 2021**.** The file provides the date, warehouse location I.D., client code, item number, and quantity.

Inventory Receipt Data

This file provides the P.O. number, client code, order quantity, order date, receipt date, and quantity received.

Order Cuts Data

The order cuts data included the retail orders that were not filled due to the unavailability of inventory in the warehouse. Data was gathered for the same time frame as the order data.

3.2 Data Preparation

Order data

The item description column and the customer reference number were used to merge the 11 CSVs into one single CSV. Furthermore, we merged the order cut data set into the merged order data to arrive at the consolidated order data for C.H. Robinson.

- **The data type of each column field -** We identified the data type of each field and observed that some data types were erroneous given the nature of the field. For example – objects were noticed in the quantity ordered field, which should have been integer-only. On looking at the line items we observed a potential shift in the data field while exporting the files into CSV. On discussing with our capstone partner, he agreed that the data entries may have been shifted in between columns, which caused the error and we decided to get rid of those line items. The cut was approximately 400,000 line items out of a total of 4000,000 line items.

- **Incorporated complete order cuts –** A preliminary test of the order fill rate metric showed a higher than expected fill rate. After data exploration, it was established that the data set was missing orders that were cut completely. We then obtained the updated order data set and ran our preliminary tests again.

- **Identifying and removing duplicates –** While running the preliminary tests on the order data set, we observed an irrational trend. The order volumes were declining from the year 2020 to 2021. 2021 should ideally have seen an increase given the reopening of the economy after the complete lockdown. We suspected erroneous duplication of order data while extracting data from their software package because, the order number, item number, and quantities were repeating in many line items. Thereafter, we received updated order data from C.H. Robinson excluding the duplicate line items.

- **Dropping erroneous values –** Item numbers being central to our analysis, we needed to make sure that the item numbers do not contain any erroneous entries such as null values and special characters. Therefore, we filtered out the null values and special characters and dropped them from a data frame.

- **Removal of non-active and partially active SKUs –** For our research, we decided to focus only on active SKUs. Therefore, we categorized the SKUs with no demand for at least 1 month as non-active/partially active SKUs. Thereafter we decided to marginalize those non-active/partially active SKUs. What remained was the list of active SKUs for 3 years.

- **Monthly aggregation –** Aggregated forecasts are more accurate than dis-aggregated forecasts. The idea is that aggregation leads to a pooling effect that will in turn lessen the variability. The peaks balance out the valleys. The coefficient of variation (CV) is commonly used to measure variability and is defined as the standard deviation over the mean (CV=σ/μ). Forecasts that are aggregated based on time (demand over a month versus over a single day) generate a much more reliable coefficient of variance for further analysis **(*MITx_MicroMasters_SCM_KeyConcepts.Pdf*, n.d.).** Therefore, we decided to aggregate the daily order data into weekly and monthly bins to reduce the variability, resulting in a forecast with relatively better root mean squared error.

3.3 Forecasting

Forecasting being central to building an inventory strategy, we created a forecasting mechanism. Furthermore, as C.H. Robinson does not operate on a forecast, we thought this would be a valuable addition to their operations. Therefore, post data cleaning and preparation, we used the order demand for the active SKUs and ran a forecast for every active SKU. We selected two forecasting algorithms to test our data with. We intended to choose the algorithm with the lowest root mean squared error (RMSE) and mean absolute percentage error (MAPE) at an SKU level. RMSE is a quick measure of how accurate the selected algorithm is with respect to forecasting actual results. MAPE is the average of the absolute percentage errors of forecasts.

3.3.1 Forecasting Algorithms Applied:

Autoregressive Integrated Moving Average (ARIMA)

ARIMA model is a combination of the differenced autoregression model and the moving average model. The autoregressive model is a time series model, that uses observed values from historical data as an input to drive prediction of the future behavior. The best predictor of future instances of a variable is the past instances of that same variable. Moving average or rolling mean is calculating a simple average for a specified period time-period out of the entire time frame of historical data. Therefore, all observations in the specified time are given the same weight in the forecast model.

In an autoregressive integrated moving average model, the future value of a variable is assumed to be a linear function of several past observations and random errors (Zhang, 2003).

The full model to the predict a variable based on an autoregressive integrated moving average is shown in Equation 1:

Equation 1

Future value of a variable y'_t based on autoregressive integrated moving average

$$y'_t = c + \emptyset_1 y'_{t-1} + \ldots + \emptyset_p y'_{t-p} + \theta_1 \varepsilon_{t-1} + \ldots + \theta_q \varepsilon_{t-q} + \varepsilon_t,$$

where, y'_t is the differenced series (it may have been differenced more than once). The “predictors” on the right-hand side include both, lagged values ($\emptyset_1 y'_{t-1} + \ldots + \emptyset_p y'_{t-p}$) and lagged errors ($\theta_1 \varepsilon_{t-1} + \ldots + \theta_q \varepsilon_{t-q} + \varepsilon_t,$).

ETS (Error, Trend, Seasonality)

ETS is an exponential smoothing time series model. This model assumes the time series includes some trend and seasonality in the data. While ARIMA assumes all past data has equal weight, ETS lowers the weight of older data exponentially and puts more weight on the most recent data.

There are two types of ETS models, one with additive errors and the other with multiplicative errors. Essentially, the additive error model defines the error as the difference between the measurement and the truth, while the multiplicative error model defines the error as the ratio between the two. Where the selection of which model to choose depends on the data set, the multiplicative model usually is a superior choice. The reason is that, the additive model tries to fit the data set to a linear function “i.e.” a straight line. Therefore, it does not capture a lot of systematic errors (Tian et al., 2013).

Simple exponential smoothing with additive errors

A general forecasting equation using simple exponential smoothing with additive errors is shown in Equation 2

Equation 2

Forecasting variable y based on simple exponential smoothing with additive errors

$$\hat{y}_{t+1/t} = l_t$$

Where,

$\hat{y}_{t+1/t}$ = Forecast for time t+1 at time t

l_t = The previous level

If we re-arrange the smoothing equation for the level, we get the "error correction" form shown in Equation 3:

Equation 3

Error correction form of the smoothing equation with additive errors

$$l_t = l_{t-1} + \alpha(y_t - l_{t-1})$$

$$= l_{t-1} + \alpha e_t$$

Where $e_t = y_t - l_{t-1} = y_t - \hat{y}_{t|t-1}$ is the residual at time t.

The training data errors lead to the adjustment of the estimated level throughout the smoothing process for $t = 1, \ldots, T$. For example, if the error at time t is negative, then $y_t < \hat{y}_{t|t-1}$ and so the level at time $t - 1$ has been over-estimated. The new level l_{t-1} adjusted downwards. The closer α is to one, the "rougher" the estimate of the level (large adjustments take place). The smaller the α, the "smoother" the level (small adjustments take place). (Hyndman & Athanasopoulos, n.d.)

Multiplicative model with level, trend, and seasonality

The multiplicative model uses α (level), b (trend), and γ (Seasonality) parameters to assign weights to more recent data as compared to old data. Figure 8 below shows the graphical representation of the components of the multiplicative ETS model. Since the value of α (level), b (trend), and γ (Seasonality) parameters can be between 0 and 1 only. In a scenario where, the parameters α, b, and γ take a value of 1, most weight is being placed on recent data. On the contrary, if they take the value of 0, most weight is being placed on old data. The abovementioned parameters act as a lever to create a point forecast that gives the best results (Hyndman & Athanasopoulos, n.d.). Equation 4 showcases the mathematical formulae to calculate the forecast, level, trend, and seasonality estimate based on a multiplicative model.

Equation 4

Forecast, level, trend, and seasonality estimate based on multiplicative model

$$y_t = (l_{t-1} + b_{t-1})s_{t-m}(1 + \varepsilon_t)$$

$$l_t = (l_{t-1} + b_{t-1})(1 + \alpha\varepsilon_t)$$

$$b_t = b_{t-1} + \beta(l_{t-1} + b_{t-1})\varepsilon_t$$

$$s_t = s_{t-m}(1 + \gamma\varepsilon_t)$$

Where,

y_t – forecast

l_t – level estimate

b_t – trend estimate

s_t – seasonality estimate

α – level smoothing parameter

β – trend smoothing parameter

γ – seasonality smoothing parameter

ε_t – error term

Figure 8

Graphical Representation of Components of Multiplicative ETS model

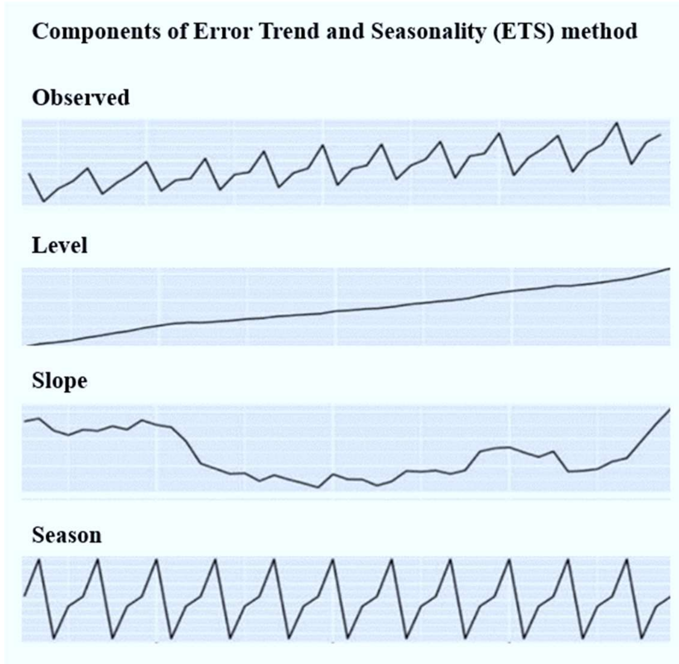

Adapted from (Hyndman & Athanasopoulos, n.d.)

3.5 Validation

To avoid overfitting our forecast, we cross-validate it. We trained and tested the prediction on a rolling basis from the year 2019 through 2020. For creating training and testing sets, we aggregated the time series data of the orders on a monthly level.

3.6 Selection of Forecast based on Key Performance Indicator (KPI)

After building a time series forecasting model, and cross-validating it, we selected a forecasting technique based on forecast accuracy. Forecast accuracy was measured by summarizing forecast errors using the key performance metrics of root mean squared error (RMSE) and mean absolute percent error (MAPE). Forecast errors (e_t) are the difference between the actual demand and the forecast for the demand at time t. Forecast errors are calculated on the test set.

Equations 5 and 6 showcase the two prominent ways to measure forecast accuracy that we reviewed, root mean squared error (RMSE) and mean absolute percentage error (MAPE).

Equation 5

Root Mean Squared Error (RMSE)

$$\text{RMSE} = \sqrt{(mean(e_t^2)}$$

Equation 6

Mean Absolute Percent Error (MAPE)

$$\text{MAPE} = mean\left(\left|100\,\frac{e_t}{y_t}\right|\right)$$

Where,

e_t = forecast errors

y_t = observed value of demand

We selected the forecasting method that returned the lowest MAPE. The reason is, that it gives a percentage error which has the advantage of being unit-free, therefore making it effective in comparing forecast performances between data sets. The MAPE minimizes the percentage of forecast errors (e_t) over the actual observed value of demand. Therefore, the lowest MAPE would represent the forecast with the highest accuracy.

3.7 Clustering

As mentioned in the literature review in section 2.8, different item categories would have different behaviors. Thus, it is important to first categorize these items before creating an inventory policy.

Therefore, we will use clustering to select SKUs that represent the same behavior.

Our overall process for clustering is the following:

- Select top 50 SKUs (out of 3,769 SKUs) representing high volumes of demand. Rank the top 50 SKUs based on the highest order quantity (measured in cases). These SKUs represent a large percentage of total business, but their fill rate is declining faster than the rest of the SKUs. Therefore, developing an inventory strategy to alleviate the fill rate of these SKUs would be of high importance to C.H. Robinson.
- Select three variables that would affect the behavior of the SKUs. These are mean absolute percent error of forecast, coefficient of variance of SKUs, and fill rate.
- Run the k-means cluster algorithm with the output of 3 clusters based on 3 metrics that best describe the demand behaviors of the SKUs.

3.8 Inventory Strategy

Once we identified the active SKUs, we obtained a forecast at an SKU level and clustered SKUs that showed similar behavior. Thereafter, we built an inventory strategy for a selection of SKUs. For building an inventory strategy, we focused on a periodic review policy.

3.8.1 Periodic Review Policy (Order-Point, Order-Up-to-Level (R, S) System):

This system, also known as a replenishment cycle system, is in common use, particularly in companies without sophisticated computer control. It is also frequently seen when items are ordered from the same supplier or require resource sharing. The control procedure is that every R unit of time ("i.e.", at each review instant) enough is ordered to raise the inventory position to the level S (Silver et al., 2017). Figure 9 below shows the periodic review control system in effect. It can be seen in the graph that after constant review periods, orders are placed to bring the inventory position to the order-up-to-level (S).

Figure 9

Illustration of a Periodic Review Policy (Order-Point, Order-Up-to-Level (R, S) System)

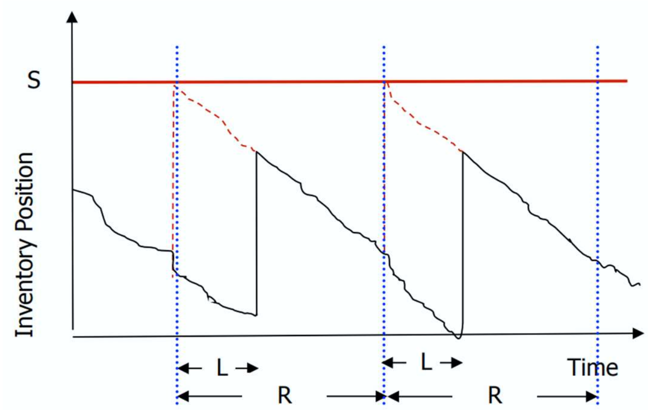

Source: (*DURING CLASS - Slides and Other Class Materials-2: SCM.260 Logistics Systems*, n.d.)[PowerPoint slides].Canvas.mit.edu

C.H. Robinson does not make decisions with respect to the quantity ordered and order frequency. However, the level of safety stock that is maintained at their distribution center impacts their fill rate performance. Therefore, for our research, we focused on the Periodic Review Policy. The reason is, that the periodic review policy enables setting a safety stock and sets the order up to level. Therefore, it will give a benchmark to our capstone partner for the optimal safety stock and order quantity to reach an order fill rate of 98%.

Determining the review interval (R)

In an (R, S) control system, a replenishment order is placed for every R unit of time; and when computing the value of S, we assume that a value of R has been predetermined (Silver et al., 2017).

C.H. Robinson does not implement the periodic review policy now. Therefore, we tested various review periods based on their business processes to get the best result.

Determining lead time:

Since C.H. Robinson does not place any orders to their customers, there is no lead time with respect to order and receipt of inventory. Therefore, based on our understanding of the business processes, we assumed that the lead time for our research is 1 month.

Determining safety factor (k) based on item fill rate:

The item fill rate (IFR) is the fraction of demand that is met with the inventory on hand. This is frequently used as a performance metric where the inventory policy is designed to minimize the cost to achieve an expected IFR of 98. If the target IFR is known then the appropriate k value can be ascertained by using the Unit Normal Loss Function, *G(k)*. The unit normal tables provide the value of *k* that corresponds to a given value of *G(k).* Equation 7 showcases the equation to calculate the value of *G(k)*.

Equation 7

Mathematical calculation to ascertain G(k)

$$G(k) = \frac{Q}{\sigma_{D_{L+R}}}(1 - IFR)$$

Where,

$G(k)$ – Unit normal loss function value for a given k

Q – Order quantity

$\sigma_{D_{L+R}}$ – Standard deviation of forecast errors over lead time and the review period

IFR – Item fill rate

(*MITx_MicroMasters_SCM_KeyConcepts.Pdf*, n.d.)

In C.H. Robinson's case, given that the target order fill rate is 98%, we have taken the item fill rate to be 98% for our research. Furthermore, in place of $\sigma_{D_{L+R}}$, we used RMSE of forecast to ascertain *G(k)*. The reason is that C.H. Robinson does not operate on a forecast and therefore it was not possible to ascertain forecast errors. Therefore, we used equation 8 to ascertain *G(k)* and thereby the safety factor (*k*):

Equation 8

Mathematical calculation we used to ascertain G(k) for our research purposes

$$G(k) \; = \frac{Q}{RMSE}(1 - IFR)$$

Where,

$G[k]$ – Unit normal loss function value for a given k

Q – Quantity ordered

$RMSE$ – Root mean squared error of the item forecast

IFR – Item fill rate

Determining the Order-Up-to-Level (S):

In periodic review policy, the key time over which protection is required is duration R + L, instead of just a replenishment lead time L. This is because once an order is placed based on the order up to level (S), it should be sufficient to cover the demand for the period of duration R+L. A stockout will occur at the end of the current cycle if the total demand in an interval of duration R + L exceeds S (Silver et al., 2017). Equation 9 shows the order up to a point (S) calculation:

Equation 9

Mathematical calculation used to ascertain order-up-to level (S)

$$S \; = \; X_{DL \, + \, R} \; + \; k.RMSE$$

Where,

X_{DL+R} : mean demand over lead time and review period

RMSE: root mean squared error of the forecast
k: safety factor
(*MITx_MicroMasters_SCM_KeyConcepts.Pdf*, n.d.)

Determining safety stock:

Once we obtain the root mean squared error (RMSE) of the forecast and safety factor (k) based on the item fill rate, we can calculate the safety stock using the equation 10 (Silver et al., 2017)

Equation 10

Mathematical calculation used to ascertain safety stock (SS)

$$SS = k.RMSE$$

Where,

RMSE: root mean squared error of the forecast
k = safety factor

3.9 Validation with Existing Inventory Levels:

After we ascertain the optimal safety stock levels to ensure an order fill rate of 98%. We will compare the optimal inventory levels for the selected SKUs to actual inventory levels. Based on this comparison we will make recommendations to C.H. Robinson to build an optimal inventory strategy that minimizes cost.

3.10 Conclusion

For our methodology, we focused on identifying SKUs that represent a significant scope of improvement in terms of the order fill rate metric for C.H. Robinson. Furthermore, we relied on the literature on inventory management to develop a strategy to increase the order fill rate for C.H. Robinson. In chapter 4, we examine the data analysis and results from the application of our methodology to the data set.

4. RESULTS AND ANALYSIS

As per the methodology described in Figure 7, we first garner insights about the existing order fill rate metric from the data received from C.H. Robinson. After data exploration, we identify the active SKUs. Thereafter, within active SKUs, we identify the top 50 SKUs that consistently represented a significant portion of the business. We also examine results from building a forecast for the SKUs using ARIMA and ETS methods. Subsequently, we perform clustering and segment the SKUs that follow similar behavior. Finally, we calculate the safety stock and order-up-to-level for a selection of SKUs to build an inventory strategy.

4.1 Data Exploration

We examine the order data to ensure that the data provided is consistent with the problem statement of the declining order fill rate metric. We group the order data to obtain an aggregate of orders placed by each customer for three years. In addition to that, we sort the sales data by customer based on high-to-low values to obtain the top five suppliers for C.H. Robinson's business. Figure 10 shows the percentage share of total revenue for the top 5 suppliers (Y axis) through the years 2019 to 2021 (X-axis).

Figure 10

Share of Total Revenue Attributable to Top 5 Suppliers

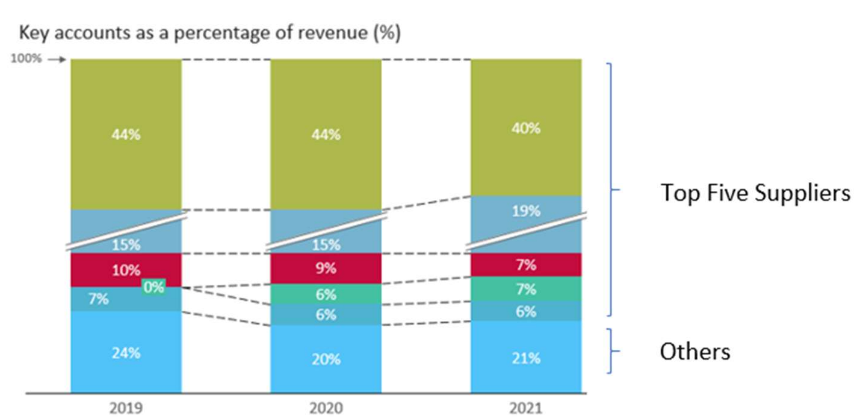

The order data received from C.H. Robinson represents sales of 13 suppliers of the company for the year 2021. As observed in Figure 8, the top five suppliers contribute 79% of C.H. Robinson's business in the year 2021, whereas the others represent 21% of its business in the year 2021.

4.2 Comparison of Sales Volume for the Top Five Suppliers

After identifying the key suppliers for C.H. Robinson's business, we observe the year-over-year sales volume growth. Table1 shows the sales volume comparison for the key suppliers over the years.

Table 1

Year-Over-Year Comparison of Sales Volume for Key Suppliers

	Y-o-Y Change in Sales Volume		
Supplier	**19 to '20**	**20 to '21**	**Fill Rate '21**
Supplier 1	50%	-15%	95%
Supplier 2	51%	22%	84%
Supplier 4	4178%	-4%	94%
Supplier 5	27%	5%	83%
Other	50%	6%	92%

From the data provided in Table 1, we observe that Suppliers 2 and 3 grew in terms of their sales volume. Supplier two has double-digit growth, both from 2019 to 2020 and from 2020 to 2021. However, the fill rate was 84% in 2021. Similarly, for supplier five, the fill rate was 83% despite increasing volumes. Overall, it is clear from examining the sales data that the order fill rate metric is decreasing. Moreover, the customer order data analysis provides an insight into which key accounts are more problematic than the others.

4.3 Active SKU Identification

Upon examining the order data, it is evident that C.H. Robinson manages 3,769 SKUs in total. While conducting preliminary data analysis, it is established that not all SKUs had consistent sales throughout the three years' worth of data provided to us. Therefore, we categorize SKUs as active, partially active, and not active SKUs based on their sales volume for every year. The SKUs that have absolutely no sales in a year are categorized as not active. Moreover, the SKUs which have no sales volume in at least one whole month of a year, are defined as partially active SKUs. Finally, the SKUs that remain are SKUs that have sales transactions every month of every year from 2019 to 2021. Therefore, we categorize those SKUs as active SKUs. After the identification of active SKUs, we reduce the scope of our analysis to these active SKUs. By doing so, we focus on addressing the problem of decreasing the fill rate metrics for SKUs that are consistently representing the bulk of C.H. Robinson's business. Figure 11 shows the segregation of SKUs into active, partially active, and not active from the year 2019 to 2021.

Figure 11

Segregation of SKUs Based on Activity for the Years 2019 through 2021

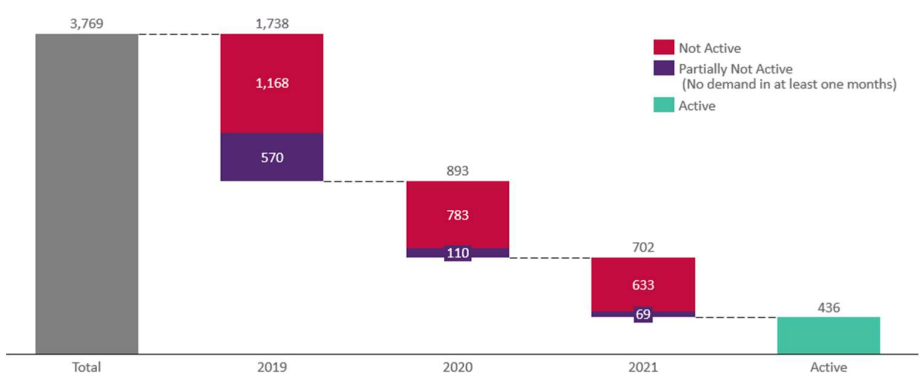

As seen in Figure 11, in the year 2019, 1168 SKUs did not have any sales transactions and are categorized as not active. There are 570 SKUs that have no demand in at least one whole month and are categorized as partially active. Therefore, a total of 1,738 inactive and partially active SKUs are eliminated. Similarly, for 2020 and 2021 a total of 893 and 702 SKUs are eliminated for being partially and/or completely inactive. After eliminating the partially and completely inactive SKUs we obtain 436 active SKUs highlighted in green. These SKUs are the ones sold consistently every month over a period of three years.

4.4 Identifying Top 50 SKUs for Analysis

We sort the 436 active SKUs from high-to-low values to obtain 50 SKUs with the highest sales volume over the period of three years. Thereafter, as shown in Figure 12, we compare the sales volume and fill rate of the top 50 SKUs with the rest, for the years 2019 through 2021.

Figure 12

Sales and fill rate comparison of Top 50 SKUs through 2019 – 2020

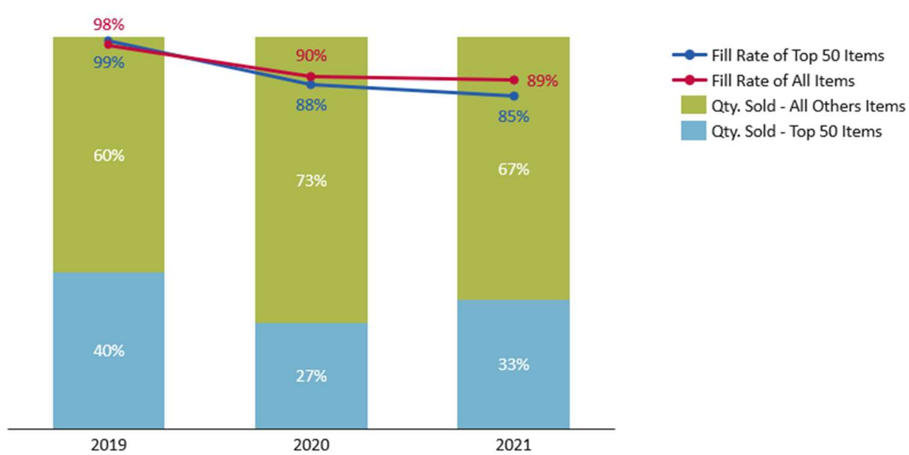

In 2019, the top 50 SKUs represent 40% of sales volume attributable to active SKUs. Despite representing such a significant share of revenue, they have a high order fill rate of 99% in the year 2019. In the year 2021, the top 50 SKUs represent a 33% share of total revenue attributable to active SKUs. Even though the percentage share declines from 40% to 33%, it is still a significant share of the business represented by 50 SKUs. Furthermore, we observe that the fill rate declines from 99% in 2019 to 85% for the top 50 SKUs. That signifies a decreasing order fill rate for SKUs that have consistently represented a significant share of the business. Therefore, we focus on the top 50 SKUs to derive an inventory strategy.

4.5 Forecasting for Top 50 Stock Keeping Units

We use the autoregressive integrated moving average (ARIMA) and error trend and seasonality (ETS) method to build a forecast for the top 50 SKUs identified in section 4.3 of the analysis and result. We build a forecast using the rolling cross-validation method reviewed in section 2.5 of the literature review for both ARIMA and ETS. Thereafter, we compare the mean absolute percentage error across the two methods. As mentioned in section 2.6 in the literature review, MAPE calculates the average of percent forecast errors. The lower the MAPE, the higher the accuracy of the forecast. Figure 13 shows the MAPE using both ARIMA and ETS methods for the top 50 SKUs.

Figure 13

MAPE of forecast using ARIMA and ETS for Top 50 SKUs

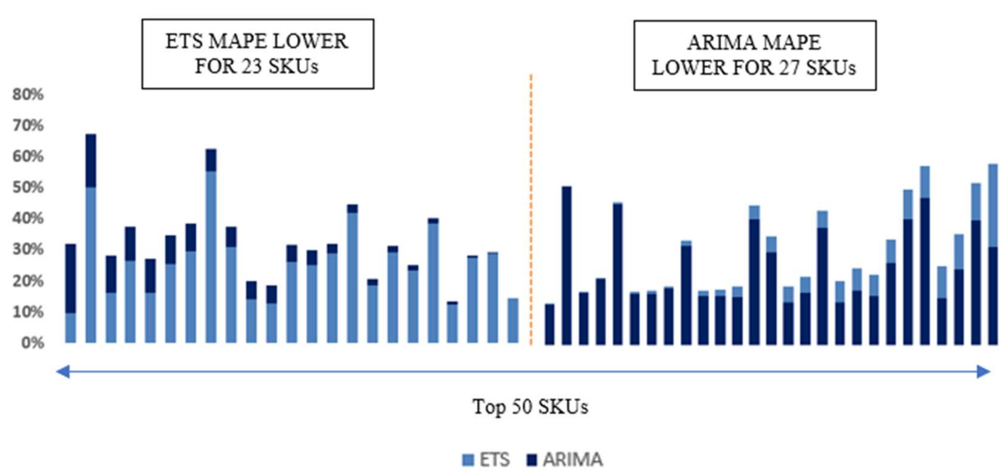

As shown in Figure 13, we observe that for 23 SKUs the ETS method produces forecasts with a lower MAPE as compared to ARIMA. Whereas, for 27 SKUs, the ARIMA method generates a lower MAPE. Finally, we select the method which produces the lower MAPE to build an inventory strategy for those products. In the case of the 23 SKUs for which ETS is selected, the average MAPE is ~26%, whereas, in the case of the 27 SKUs for which ARIMA is selected, the average MAPE is ~24%.

4.6 Clustering Top 50 Stock Keeping Units

As mentioned in section 2.8 of the literature review, certain SKUs would have different demand patterns. To identify which products follow a similar characteristic of demand, we cluster the top 50 SKUs based on three key variables. We select the following three variables to cluster SKUs:

1. **Coefficient of variation ($\boldsymbol{\sigma/\mu}$)**: The coefficient of variation is calculated by dividing the standard deviation of demand by average demand. It is a metric used to calculate the relative variability in demand. The higher the coefficient of variance, the more volatile the demand for a given SKU is.
2. **Order fill rate** ($\frac{\boldsymbol{Qty.\ Shipped}}{\boldsymbol{Qty.\ Ordered}}$): For all the top 50 SKUs, we calculate the aggregate order fill rate over the period of three years. To do so, we divide the total quantity shipped by the total quantity ordered to obtain the order fill rate for an SKU.
3. **Mean absolute percent error ($\boldsymbol{\frac{1}{n}\Sigma\frac{|e\_t|}{d\_t}}$):** MAPE is a forecast key performance indicator that indicates the forecast accuracy of an SKU. The lower the MAPE, the more accurate a forecast is.

Thereafter, we cluster the top 50 SKUs using the k-means clustering method. Figure 14 shows the three clusters of SKUs we obtained.

Figure 14

Top 50 SKUs divided into three Clusters

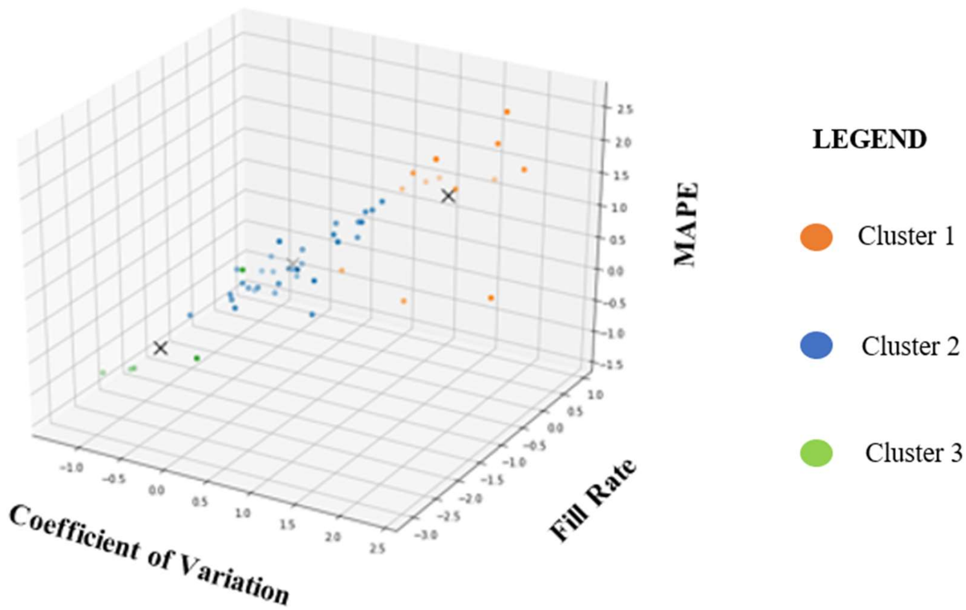

From analyzing Figure 14, we observe the three clusters based on the coefficient of variance, order fill rate, and mean absolute percentage error. The orange, blue and green dots represent each of the top 50 SKUs divided into three clusters. However, since the graph obtained from running k-means clustering was in three dimensions, we could not closely examine each variable's impact on clustering the SKUs. Therefore, we break down Figure 14 into three Figures from 15 through 17, each showing the impact of two variables in clustering the SKUs at a time.

Fill Rate vs. Variability

We mapped the top 50 SKUs with a coefficient of variance on the x-axis and order fill rate on the y-axis. Given their definitions, an SKU with a higher variability was more likely to have a lower order fill rate. Figure 15 shows the impact of the coefficient of variance and order fill rate on the three clusters.

Figure 15

Impact of Coefficient of Variance and Order Fill Rate on the three Clusters

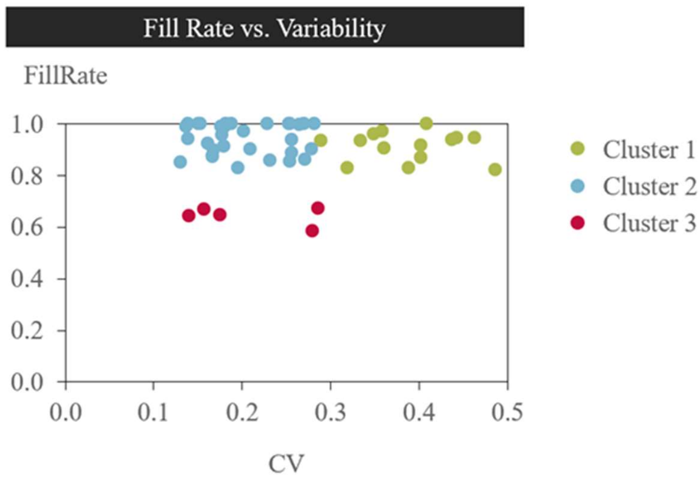

For the SKUs identified in Cluster-1, though the variability is on the higher side, the order fill rate metric is relatively high, in the 80% to 100% range. In the case of Cluster-2, the CV is low in the range of 10% to 30%. Therefore, it is justified to have a high fill rate in the range of 80% to 100% for those SKUs. Finally, the SKUs in Cluster-3 had low variability in the range of 15% to 30%. However, the order fill rate is on the lower end, between 60% to 70%. Therefore, the SKUs in Cluster-3 represented the cluster with the maximum opportunity for improvement in performance in the order fill rate metric.

MAPE vs. Variability

After analyzing the impact of variability and fill rate on the clustering of SKUs, we examine the impact of MAPE and variability on the cluster of SKUs. Theoretically, an SKU with higher variability must have a higher MAPE in building a forecast. Figure 16 shows the impact of MAPE and Variability in clustering the SKUs.

Figure 16

Impact of MAPE and Variability in Clustering the SKUs

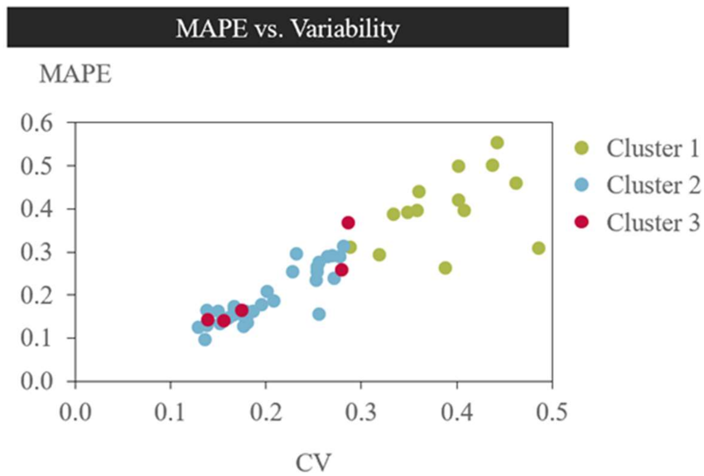

On examining the graph, it was evident that the SKUs with higher variability had a higher MAPE across clusters. We observed that there was a linear relation between CV and MAPE throughout the three clusters, i.e., as the CV increased on the x-axis, the MAPE increased on the y-axis. Therefore, there is a positive correlation between CV and MAPE observed in Figure 16.

MAPE vs. Fill Rate

Finally, we examine the effect of MAPE and order fill rate in clustering the top 50 SKUs. Theoretically, in the case of MAPE and order fill rate, the higher the MAPE, the lower should be the order fill rate. Figure 17 illustrates the impact of MAPE and order fill rate on the three clusters.

Figure 17

Impact of MAPE and Order Fill Rate in Clustering SKUs

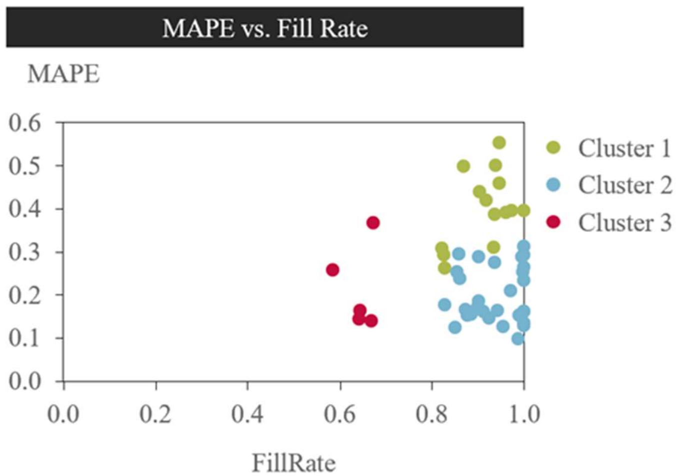

For Cluster-2, the fill rate is in the higher range of 80% - 100%, and the MAPE on the lower end is between 10% and 30%. In the case of Cluster-1, the fill rate is in the higher bracket of 80% - 100% despite a higher MAPE range of 30% to 60%. Finally, for SKUs in Cluster-3, the order fill rate is in the lower range of 55% to 70% even with a relatively low MAPE range of 20% to 40%. Therefore, the impact of MAPE and order fill rate is not consistent with the theoretical relation shown by order fill rate and MAPE.

4.7 Building a Periodic Review Inventory Control System for Active Stock Keeping Units

As mentioned in section 3.8, building an inventory strategy of the methodology, after identifying the active stock keeping units, forecasting, and clustering, we designed an appropriate inventory control system based on periodic review for each cluster.

According to the periodic review inventory control system described in 3.8.1, we identified:

Review period - How often the inventory status should be determined

Based on common industry practices in retail, we took a set of 4 review intervals. The 4 intervals we selected were 1-week, 2-weeks, 3-weeks, and 4-weeks. The review intervals are based on the trade-off between labor cost to monitor inventory levels, and inventory holding determined by the size of order quantity. A shorter review period of 1-week entails a higher labor cost and a lower inventory holding cost, as inventory levels need to be monitored more frequently and order size is small. On the other hand, a longer review interval of 4-weeks leads to a lower labor cost and higher inventory holding cost, as inventory needs to be monitored for a lesser time and the size of the order is large. We analyzed the impact of each of the four review periods on the safety stock and order-up-to level to give the best possible recommendation to the sponsor company.

Lead Time – Time between the date suppliers receives the order and the date of receipt of inventory at the C.H. Robinson distribution center

Based on industry practices and discussion with the sponsor company, we assume the lead time to be 1 month.

Safety factor (k) based on item fill rate of 98%

Given that the target order fill rate is 98%, we set our target item fill rate (IFR) for the periodic review control system to be 98%. Therefore, we use equation 8 shown in section 3.8.1 of the methodology chapter to obtain the unit normal loss function value (G[k]) corresponding to the IFR of 98%. After obtaining the G[k] corresponding to an item fill rate of 98%, we use the unit normal tables to arrive at the safety factor value (k) to calculate appropriate inventory levels as per the periodic review control system. We obtain a safety factor (k) for each of the top 50 active SKUs for our analysis using the unit normal tables.

Safety stock level – Inventory level required to buffer against uncertainty in demand

We use equation 10 mentioned in part 3.8.1 in methodology, the periodic review control system to calculate the safety stock levels. The root mean squared error (RMSE) is obtained while building a forecast for each SKU as mentioned in section 4.5. In addition to the RMSE, we use the safety factor (k) as derived from equation 7 in section 3.8.1 for each SKU to calculate the optimum safety stock level for each cluster. Figure 18 through 20 illustrates the safety stock levels across the three clusters.

Figure 18

Safety stock levels (Cases in '000s) calculated for Cluster-1 across review periods based on RMSE and safety factor (k)

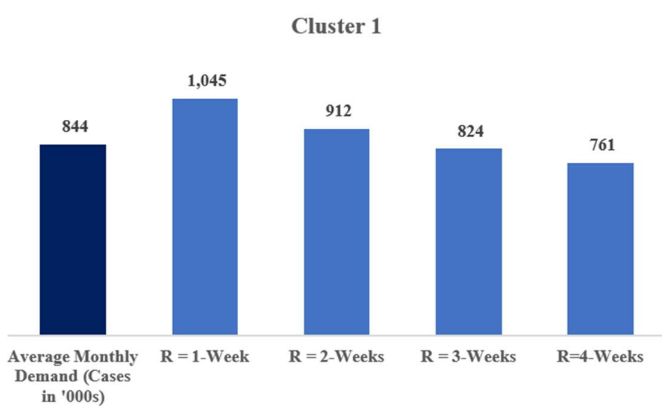

As per Figure 18, a safety stock of over 1 million cases needs to be maintained if the review period is 1 week, whereas, a safety stock of 761,000 cases if the review period is 4-weeks. There is a decreasing trend in the quantity of safety stock required to be maintained as the review period increases from 1-week to 4-weeks.

Figure 19

Safety stock levels (Cases in '000s) calculated for Cluster 2 across review periods based on RMSE and safety factor (k)

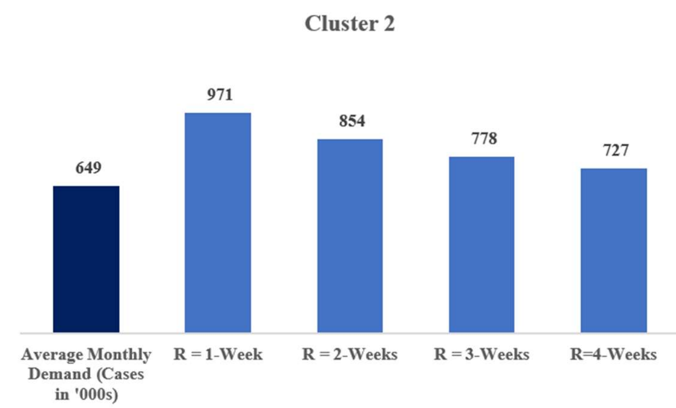

As per Figure 19, a safety stock of 971,000 cases needs to be maintained if the review period is 1 week, whereas, a safety stock of 727,000 cases if the review period is 4-weeks. There is a decreasing trend in the quantity of safety stock required to be maintained as the review period increases from 1-week to 4-weeks, like what was observed in Figure 18 for Cluster-1.

Figure 20

Safety stock levels (Cases in '000s) calculated for Cluster 3 across review periods based on RMSE and safety factor (k)

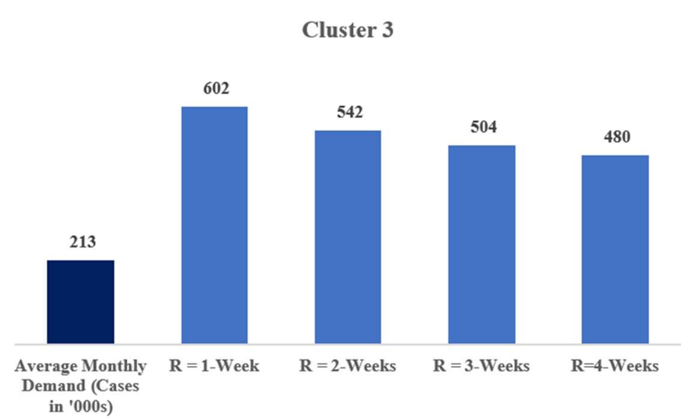

As per Figure 20, a safety stock of 602,000 cases needs to be maintained if the review period is 1 week, whereas, a safety stock of 480,000 cases if the review period is 4-weeks. There is a decreasing trend in the quantity of safety stock required to be maintained as the review period increases from 1-week to 4-weeks, like what was observed for Cluster-1 & 2 in Figures 18 and 19.

Order-up-to-level (*S*) – Inventory level to buffer against the demand between two review periods and lead time.

In the periodic review control system, since an order is placed after every constant review interval (R), the order quantity needs to be enough to buffer against the demand through the review period. Therefore, we used equation 9 shown in section 3.8.1 of the methodology to derive the order-up-to level for SKUs at a cluster level.

Figures 21 through 23 provide the order-up-to-level for SKUs at a cluster level along with the average monthly demand.

Figure 21

Average monthly demand and order-up-to level for SKUs in Cluster 1 in cases in 000's

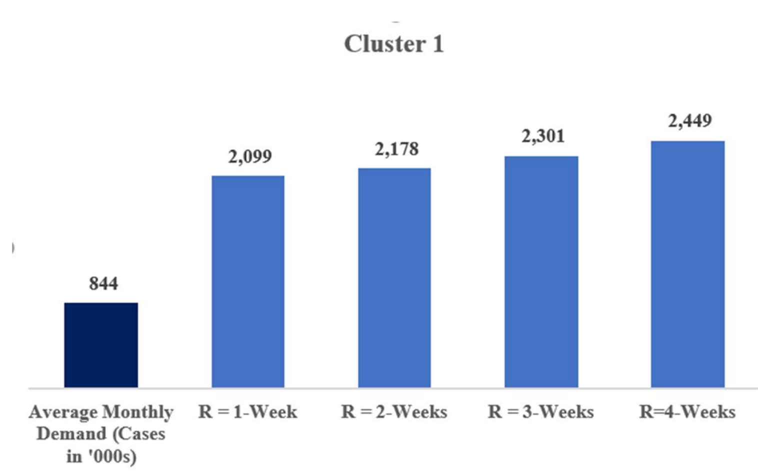

As per Figure 21, when the review period is of 1-week, the order-up-to level (*S*) of inventory is 2.09 Million cases that need to be maintained for Cluster-1. Therefore, whatever the inventory on hand at the time of review, an order must be placed to the suppliers to ensure inventory goes up to the level of 2.09 million cases. By doing so, C.H. Robinson will be able

Figure 22

Average Monthly Demand and Order-Up-To Level for SKUs in Cluster-2 in cases in 000's

to ascertain the number of cases to be received from their suppliers to attain an order fill rate of 98%.

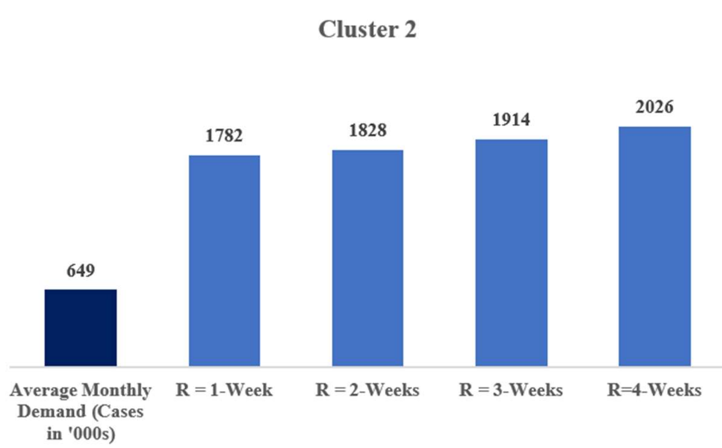

As per Figure 22, in the case of a review period of 1-week, an order-up-to level of inventory of 1.78 Million cases needs to be maintained for Cluster-2. Therefore, whatever the inventory on hand at the time of review, an order must be placed to the suppliers to ensure inventory goes up to the level of 1.78 million cases. By doing so, C.H. Robinson will be able to ascertain the number of cases to be received from their suppliers to attain an order fill rate of 98%.

Figure 23

Average monthly demand and order-up-level for SKUs in Cluster 3 in cases in 000's

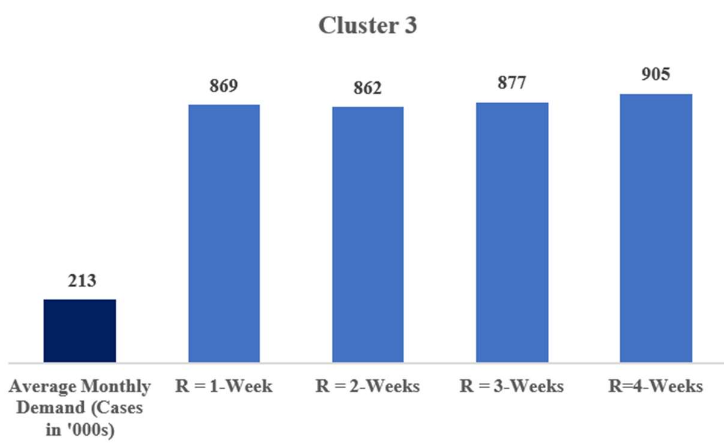

As per Figure 23, when the review period is 1-week, the order-up-to level (S) of inventory is 869,000 cases that need to be maintained for Cluster-1. Therefore, whatever the inventory on hand at the time of review, an order must be placed to the suppliers to ensure inventory goes up to the level of 869,000 cases. By doing so, C.H. Robinson will be able to ascertain the number of cases to be received from their suppliers to attain an order fill rate of 98%.

After exploring the data, it is evident that the top 50 active SKUs represent the bulk of the problem of declining order fill rate. Furthermore, on comparing forecast accuracy between ARIMA and ETS, we realize different forecasting methods are suitable for different SKUs. K-means clustering enabled segmenting of the top 50 SKUs into three clusters based on variability in demand, forecast accuracy, and fill rate. Lastly, on the application of a periodic review control system to the top 50 items, we obtained the safety stock and order-up-to levels to design an inventory strategy for the top 50 stock-keeping units.

5. DISCUSSION

5.1 Business analysis, forecasting, and Clustering of top 50 SKUs

Based on our data exploration in part 4.1, the difference in the fill rate performance of different SKUs was evident. According to the analysis illustrated in Figure 8, the active SKUs that consistently represented the bulk of C.H. Robinson's business had order fill rates that declined faster than the rest of the SKU portfolio. Therefore, building an inventory strategy for active SKUs would be efficient to implement and bring significant improvement to the order fill rate.

Building a Forecast

As mentioned in section 2.4 of the literature review, building a forecast is central to designing an inventory policy. We used ARIMA and ETS techniques to build a forecast for active SKUs. After calculating the forecast for the active stock keeping unit, we examined the forecast error, which describes the accuracy of a forecast. It is the difference between the forecast and observed values of demand in the past. When it came to selecting a forecasting method for an SKU between ARIMA and ETS, we chose the method with a lower mean absolute percent error. As per section 2.6 in the literature review, the mean absolute percent error metric is ideal for comparing two forecasting methods, as it gives out the error in interpretable terms. According to Figure 13, it is evident that for 23 SKUs ETS method gave more accurate forecasts and for 27 SKUs ARIMA method was more accurate. The root mean squared error is another renowned forecast error metric that we used, in deciding the level of safety stock a company needs to maintain to buffer against uncertainties in demand. The root mean squared error was 44% of the size of actual demand for active SKUs. This signified that the forecast was deviant by 44% of the demand size, and a prominent reason is the volatility in the demand patterns from 2019 through 2021.

Segmentation of Top Selling Stock Keeping Units

We clustered the active SKUs that showed similar demand and order fill rate characteristics to build an inventory strategy. We chose coefficient of variation (CV), order fill rate, and mean absolute percent error (MAPE) as the three key variables based on which we clustered SKUs. The three variables were chosen because they describe the variability in demand, fill rate, and accuracy of forecast respectively, all central to addressing the task of increasing the order fill rate to 98%. We used k-means clustering to identify three clusters based on the key variables. The clusters are illustrated in Figure 14, however, to understand the impact of each variable in clustering the SKUs, we plotted three two-dimensional graphs based on the key variables.

Mean Absolute Percent Error (MAPE) vs. Coefficient of Variation (CV)

When we compared MAPE and CV in Figure 16, the results showed expected behavior; that is, SKUs with a higher CV had a higher MAPE. Moreover, clusters were not easily identifiable based on Figure 16. On the other hand, Figures 15 and 17 were visually much clearer in displaying the top-50 SKUs in three clusters.

Order Fill Rate vs. Coefficient of Variation (CV)

According to Figure 15, on comparing order fill rate and coefficient of variation, we observed SKUs in Cluster-1 had a high fill rate despite a high CV. This implied good in-stock performance despite high volatility in demand. Therefore, SKUs in Cluster-1 already display effective inventory management. The SKUs in Cluster-2 have low variability and therefore, the relatively higher order fill rate performance is justified. The SKUs in Cluster-3 have a low CV, which implies low volatility in demand, nonetheless, their order fill rate metric is very low. This implies poor performance in terms of the in-stock availability of those SKUs. We suggest that C.H. Robinson should have a discussion with the suppliers of the SKUs in Cluster-

3. The reason is that the SKUs are active and in high demand, therefore ensuring proper availability of inventory would increase the order fill rate performance.

Mean Absolute Percent Error (MAPE) vs. Order Fill Rate

Figure 17 shows the impact of clustering-based MAPE and order fill rate. However, it conveys a similar message as Figure 15. The SKUs in Cluster-1 represent a higher forecast error and irrespective of that they maintain a high order fill rate, signifying effective inventory management. The SKUs in Cluster-2 have a lower forecast error, consequentially a high order fill rate. Lastly, the SKUs in Cluster-3 have a low fill rate despite a lower MAPE. A lower MAPE signifies higher forecast accuracy, which should lead to effective inventory management. However, in the case of Cluster-3 SKUs, the order fill rate is low despite a low MAPE. Therefore, the analysis based on clustering warrants a deeper look into the inventory management of Cluster-3 SKUs. Even though many SKUs in Cluster-1 and 2 are below the target fill rate metric of 98%, increasing their fill rate would be relatively easier. However, in the case of SKUs in Cluster-3, the management at C.H. Robinson will have to work closely with the suppliers to ensure an order fill rate of 98%.

5.2 Inventory Strategy

We applied a periodic review inventory control system to the data. Based on the periodic review system we have identified the optimal inventory strategy for C.H. Robinson to achieve a 98% order fill rate metric. Thereafter, identified the optimal review period, level of safety stock, and order-up-to level of inventory which determines the order size:

Review Period and Corresponding Inventory Levels

As mentioned in part 4.7, selecting a review period is based on the trade-off between the labor cost involved in monitoring inventory levels and inventory holding cost. Since assigning the review period to SKUs is highly subjective, for the three clusters identified in part 4.6, we examined the average inventory obtained from applying the periodic review system. Table 2 provides the contribution to business revenue for each cluster and Figure 24 illustrates the average inventory levels with different review periods across the clusters.

Table 2

Contribution to Revenue for each Cluster

Cluster #	No. of SKUs	% of Active SKU Sales
1	13	49%
2	31	38%
3	5	12%
		100%

Figure 24

Average Monthly Inventory Levels (Cases in '000s) for Different Review Periods Across Clusters

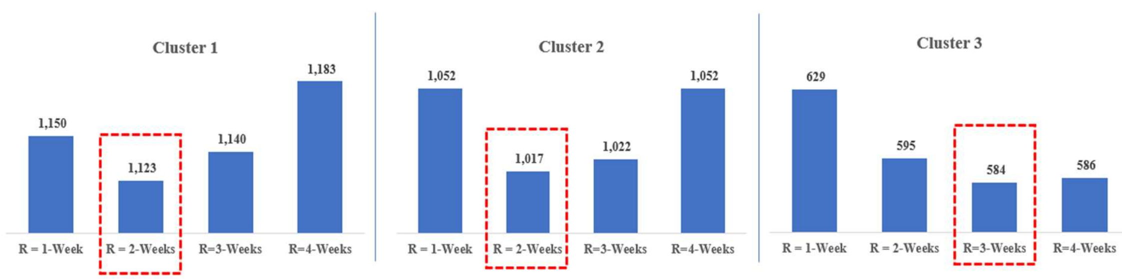

As illustrated in Table 2, Cluster 1 has 13 SKUs which account for 49% of the revenue of the top Fifty SKUs. Furthermore, on reviewing Figure 24, it is learned that a review period of 2 weeks leads to the least amount of average inventory to be held at C.H. Robinson, i.e., 1.1 million cases. The lower the average inventory levels, the lesser the holding cost incurred by suppliers of C.H. Robinson. Hence, we recommend selecting a bi-weekly review model, despite the higher cost incurred in planning for such frequent intervals. Moreover, Cluster 1 accounts for 49% of the revenue brought in by the top 50 SKUs. Therefore, incurring the cost that comes with frequent monitoring of inventory position is warranted.

In the case of Cluster 2, 31 SKUs account for 38% of the revenue brought in by the top 50 items, as illustrated in table 2. In addition to that, Figure 24 shows that a review period of 2 weeks results in the lowest amount of average inventory being held at over 1 million cases. Hence, we recommend a bi-weekly review period for Cluster 2 as well. The benefits would include lower holding cost of SKUs for suppliers that account for 38% of top 50 item revenue. Like in the case of Cluster 1, the benefits from saving money in inventory holding costs would justify the higher cost incurred in frequently monitoring inventory levels.

Lastly, for Cluster 3, 5 SKUs represent 12% of the revenue brought in by the top 50 items as per table 2. Moreover, Figure 24 shows that a 3-week review period leads to the least amount of average inventory levels the suppliers need to maintain at C.H. Robinson. Therefore, we recommend that a 3-week review period be adopted for Cluster 3. It will also result in lower inventory planning costs due to the less frequent review periods. Table 3 illustrates the summary of review periods for each Cluster and the corresponding level of safety stock and order up-to-level that C.H. Robinson should maintain to ensure an order fill rate of 98%

Table 3

Summary of Review Period Selection for each Cluster and Corresponding Inventory Levels

				Cases in '000s		
Cluster #	**No. of SKUs**	**% of Active SKU Sales**	**Review Period**	**Average Inventory Levels**	**Safety Stock Level**	**Order-up-to-level**
1	13	49%	2	1123	912	2,178
2	31	38%	2	1017	854	1,828
3	5	12%	3	584	504	877
Total	**49**	**100%**	**7**	**2724**	**2270**	**4,882**

5.3 Validation with Current Inventory Levels at C.H. Robinson

After forming the inventory strategy as illustrated in Table 3, we compared our recommended inventory levels with actual inventory levels maintained at C.H. Robinson. Figure 25 illustrates the comparison between the suggested inventory levels and the current inventory levels at C.H. Robinson.

Figure 25

Current Inventory and Suggested Average Inventory (Cases) vs. Fill Rate for Optimal Review Periods

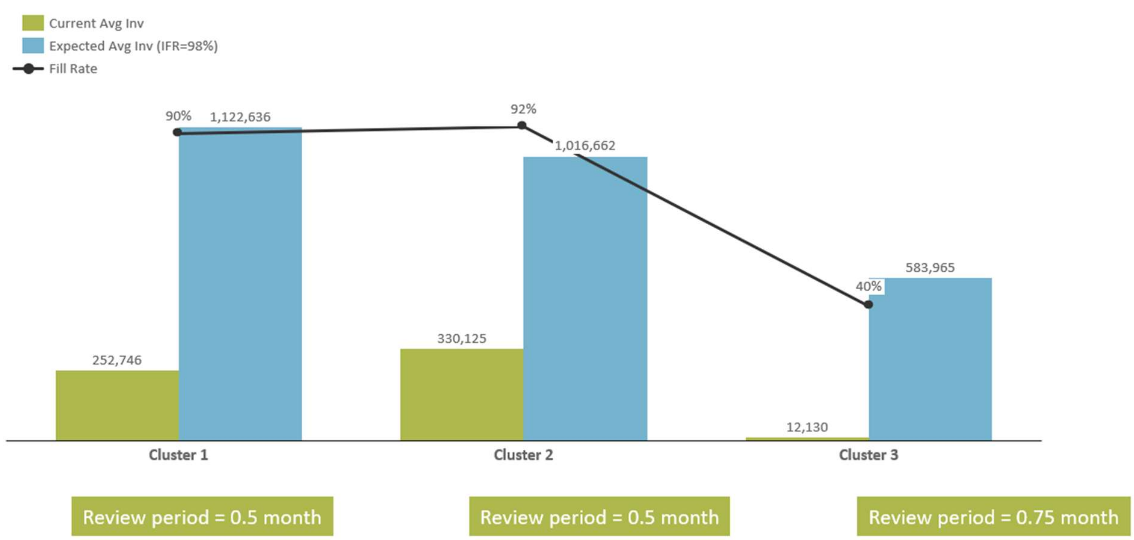

As illustrated in Figure 25, there are 252,000 cases of Cluster-1 SKUs currently maintained at C.H. Robinson's DC, which corresponds to a 90% fill rate performance. However, the recommended level of average inventory to achieve a fill rate of 98% is 1.1 million cases. That is a 444% increase in inventory levels to attain an 8% improvement on the fill rate metric.

In the case of Cluster-2, there is a recommended 402% increase in average inventory levels from 330,000 cases currently maintained at the DC, to 1.02 million cases. Therefore, by almost quadrupling the amount of inventory maintained the fill rate metric will improve from 92% to 98%.

Finally, for Cluster-3, the current order fill rate metric is at 40% by maintaining inventory levels of 12,000 cases in the DC. To increase the fill rate metric performance to 98%, the average inventory levels will need to be increased by 4,814% and maintain an average of 583,000 cases at the DC.

For all the clusters, there must be a significant increase in average inventory to attain an increase in order fill rate performance. Higher inventory levels are not desirable as they lead to increased capital commitment and holding costs. Therefore, we compared inventory levels from 95% to 98% for the selected review period for each cluster to make the best financial decision. The comparison enabled the identification of the incremental increase in inventory levels for every 1% in fill rate%. Therefore, the fill rate target leading to the best overall cost can be identified given the trade-off of inventory holding and penalization costs.

5.4 Average Inventory Levels Corresponding to Different Fill Rates from 95% through 98%

After designing the inventory strategy, and validating it with the current inventory levels, we examined the increase in average inventory levels for every increase in fill rate percentage, from 95% to 98%. We performed this analysis for all the clusters but restricted the scope to the optimal review interval selected for each cluster. Figure 26 illustrates the average inventory in cases corresponding to fill rates from 95% to 98%.

Figure 26

Average Inventory in Cases Corresponding to Fill Rates from 95% to 98%

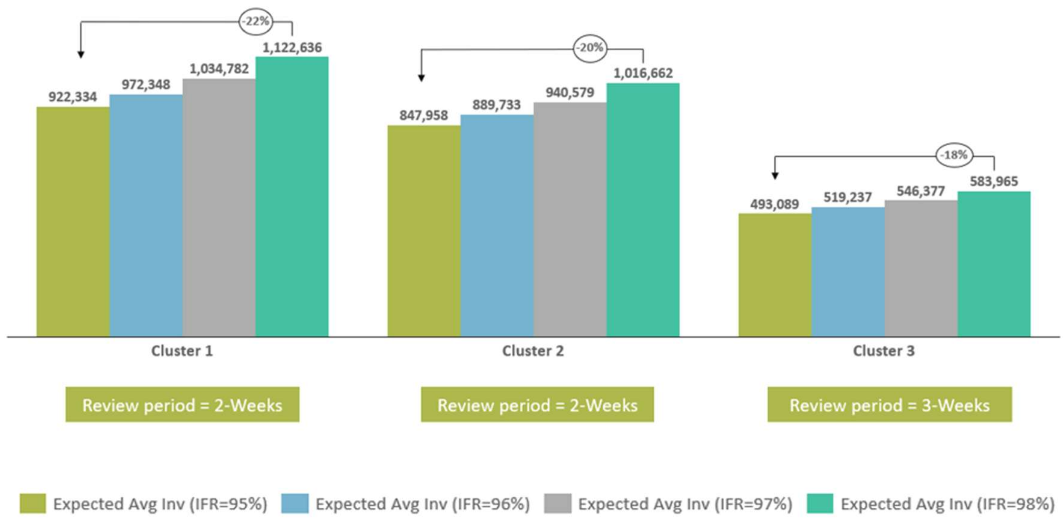

5.5 Inventory Level Recommendation based on Inventory Holding Cost vs. Penalization Cost Trade-Off

Figure 26 illustrates the average inventory levels (Cases) that suppliers need to maintain at C.H. Robinson DC for each cluster. Reviewing the graphs for all the clusters shows that it takes an increase in inventory levels by 18% - 22% to go from achieving 95% in the order fill rate metric to 98%. That is a substantial amount of money blocked in maintaining higher levels of inventory to gain 3% on the fill rate performance metric. Our suggestion to the C.H. Robinson team is that they should analyze the trade-off in money saved from holding less inventory by incurring OTIF penalties to some extent. To elucidate, if the money saved from holding inventory to ensure a 95% fill rate is significantly higher than the penalties charged by retailers from falling short of the 98% fill rate target, they should consider maintaining average inventory corresponding to a 95% order fill rate performance.

Similarly, the trade-off can be evaluated for any fill rate percentage other than 95%. Figure 26 can be used to arrive at the percentage that gives the best outcome on the inventory holding cost vs. penalty trade-off. Furthermore, depending on the dynamics of the trade-off, different order fill rate targets can be chosen for different clusters to maximize returns to C.H. Robinson's suppliers.

6. CONCLUSION

In this capstone, we focused on identifying the stock-keeping units (SKUs) that represented a continued and significant share of the overall business. As a result, we were able to identify 50 SKUs out of the total 3,769 SKUs managed by the sponsor company. Narrowing down focus like this helps in managing voluminous SKU portfolios effectively. Thereafter, we used techniques such as ARIMA and ETS to build a forecast for each of the 50 SKUs. Forecasting, using historical demand enabled us to understand the volatility in demand for each SKU through the years 2019 to 2021, which is central to inventory management decisions. The volatility in historical demand, represented by the RMSE is central to evaluating the cases of inventory to be held while designing the inventory strategy. Since every SKU has a different forecast accuracy, fill rate, and variability in historical demand, we segmented the top-50 SKUs based on these characteristics. Segmenting these SKUs enabled us to recommend 3 different inventory strategies for the 3 clusters that were more suited to their historical demand patterns. We used the k-means clustering technique to identify the 3 clusters based on the characteristics mentioned above. Finally, we used a periodic review inventory control system to build an inventory strategy for each cluster. Using the periodic review control system, we were able to answer the key research question and find the optimal review period, order-up-to level, and safety stock to achieve an order fill rate of 98%.

The goal of the project was to answer the research question so that our sponsor company can save penalty charges arising from missing 98% order fill rate targets. With the help of our advisors, we were able to not only design a strategy that answers the research question but also gives the sponsor company the flexibility of managing different SKUs differently based on their historical demand patterns. Such a dynamic and SKU-specific strategy should enable C.H. Robinson to stay ahead in the competitive retail consolidation landscape.

6.1 Managerial Insights

While examining the data to build the inventory strategy, we observed some patterns that explain the falling order fill rate metric through the years 2019 to 2021.

Order and delivery patterns

We observed that there is no consistent pattern according to which C.H. Robinson receives inventory from the suppliers. The receiving patterns are erratic and most probably because of the supply constraints caused by global supply chain challenges to the suppliers.

Quantity of inventory received by suppliers

The quantity of inventory received from suppliers is either too much or too little. It was observed that, whenever an item was received in huge quantities by a supplier, the demand slowed down. Therefore, leading to the piling up of inventory at the C.H. Robinson distribution center. To enable effective planning for inventory, the research suggests a forecast based on historical demand data, thereby avoiding piling up inventory.

Bullwhip effect caused by the shortage

Whenever a retailer's order was delivered short of the quantity ordered, the following order would be for an even bigger quantity. The reason is demand remained unfulfilled from the previous period, leading to the cumulation of order size. This pattern was observed to continue as long as there was not enough inventory at the C.H. Robinson DC. Therefore, a vicious cycle was created of orders remaining unfulfilled by a higher percentage every time an order was shorted. Therefore, the order fill rate performance was 89% towards 2021, because of the bullwhip effect that was created by orders that were shorted.

6.2 Limitations of the study and key assumptions

It is important to address the limitations in our approach that may hamper the accuracy of our inventory strategy. Furthermore, while conducting our research, a few assumptions were made which are also highlighted in this section.

- **Records of forecasts not maintained by C.H. Robinson suppliers**

Theoretically, while ascertaining safety stock and order-up-to level as per the periodic review control system, RMSE of the forecast errors is used. Forecast error (e_t) is the residual between the forecast and actual demand historically observed at time t. However, since the suppliers of our sponsor company did not maintain a record of the forecast, we built the forecast using historical demand and used the same historical demand data to calculate the RMSE. The method we used to obtain RMSE is a close approximate, and the next best alternative to the theoretical convention for our research purposes.

- **Supplier delivery lead time assumed to be 1 month**

C.H. Robinson as a consolidator does not place orders with their suppliers. The suppliers oversee inventory planning and just send inventory to the consolidator's DC as per their demand plan and production schedule. Hence, there is not a consistent supplier lead time estimate, which is key to managing inventory as per the periodic review control system. Therefore, we assumed a supplier delivery lead time of 1 month for all the SKUs based on general retail industry practices and discussion with the C.H. Robinson team.

- **Used root mean squared error of demand data to ascertain safety stock and order-up-to level, instead of $\boldsymbol{\sigma_{(L+R)}}$ of forecast errors**

We did not have the data for residual errors based on historical forecasts and actuals. Therefore, we chose to use the root mean squared error of demand data rather than forecast error to calculate safety stock and order-up-to level as per the periodic review control system.

- **Assumptions to ascertain safety factor (k), using the unit normal loss function (G[k]) equation**

We have ascertained the safety factor (k) using the unit normal loss function (G[k]) equation 8, as previously mentioned in section 3.8.1 of the methodology chapter.

Traditionally, the mathematical calculation used to ascertain G(k) for the research purposes, expected units short is used in place of RMSE. However, in the absence of data required to calculate expected units short, we have used the RMSE of demand data. The RMSE of demand signifies the volatility in demand leading to shorting of orders hence being a close approximate for E[US]. Furthermore, Q is based on the economic order quantity (EOQ)**.** However, since we did not have access to the ordering and holding cost, we used average demand over the review period to ascertain the order quantity instead of the EOQ equation.

The implication of the Research in Retail Industry

Given that the retail supply chain operations are incredibly complex, our research focuses on two key problems, i.e., penalization costs due to stocking out on items and excess inventory. Both issues lead to enormous monetary loss and require an effective inventory management strategy to obtain a solution. This research paper focuses on building an inventory strategy such that, suppliers and retail consolidators can meet order fill rate targets and minimize penalization costs. Furthermore, it also explains the trade-off involved in carrying excess inventory to achieve a high fill rate and guides management to ascertain the optimal inventory levels for their organizations, given the trade-off involved.

REFERENCES

Bower, P. (2021). Improving OTIF Metrics with Supply Chain & S&OP Best Practices. *Journal of Business Forecasting*, *40*(1), 20–28.

Chawla, A., Singh, A., Lamba, A., Gangwani, N., & Soni, U. (2019). Demand Forecasting Using Artificial Neural Networks—A Case Study of American Retail Corporation. In H. Malik, S. Srivastava, Y. R. Sood, & A. Ahmad (Eds.), *Applications of Artificial Intelligence Techniques in Engineering* (pp. 79–89). Springer. https://doi.org/10.1007/978-981-13-1822-1_8

DURING CLASS - Slides and other class materials-2: SCM.260 Logistics Systems. (n.d.). Retrieved May 5, 2022, from https://canvas.mit.edu/courses/11566/pages/during-class-slides-and-other-class-materials-2-2?module_item_id=380093

Fattah, J., Ezzine, L., Aman, Z., El Moussami, H., & Lachhab, A. (2018). Forecasting of demand using ARIMA model. *International Journal of Engineering Business Management*, *10*, 1847979018808673. https://doi.org/10.1177/1847979018808673

Garbade, M. J. (2018, September 12). *Understanding K-means Clustering in Machine Learning | by Dr. Michael J. Garbade | Towards Data Science*. https://towardsdatascience.com/understanding-k-means-clustering-in-machine-learning-6a6e67336aa1

Hyndman, R. J., & Athanasopoulos, G. (n.d.). *Forecasting: Principles and Practice*. 505.

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

Jackson, I., Avdeikins, A., & Tolujevs, J. (2019). Unsupervised Learning-Based Stock Keeping Units Segmentation. In I. Kabashkin, I. Yatskiv (Jackiva), & O. Prentkovskis (Eds.), *Reliability and Statistics in Transportation and Communication* (pp. 603–612). Springer International Publishing. https://doi.org/10.1007/978-3-030-12450-2_58

Maglogiannis, I. G. (2007). *Emerging artificial intelligence applications in computer engineering: Real word AI systems with applications in eHealth, HCI, information retrieval, and pervasive technologies*. IOS Press.

MITx_MicroMasters_SCM_KeyConcepts.pdf. (n.d.).

New Research Report: Retailers Lose $1.75 Trillion in Revenue Worldwide Due to Overstocks, Out-of-Stocks, and Returns. (2015, May 6). https://www.businesswire.com/news/home/20150506005233/en/New-Research-Report-Retailers-Lose-1.75-Trillion-in-Revenue-Worldwide-Due-to-Overstocks-Out-of-Stocks-and-Returns

Ramanathan, R. (2006). ABC inventory classification with multiple-criteria using weighted linear optimization. *Computers & Operations Research*, *33*(3), 695–700. https://doi.org/10.1016/j.cor.2004.07.014

Silver, E. A., Pyke, D. F., & Thomas, D. J. (2016). *Inventory and Production Management in Supply Chains*. Taylor & Francis Group. http://ebookcentral.proquest.com/lib/mit/detail.action?docID=4771754

Silver, E. A., Pyke, D. F., & Thomas, D. J. (2017). *Inventory and Production Management in Supply Chains* (Fourth edition). CRC Press. https://doi.org/10.1201/9781315374406

The supply-chain solution to aid retailers' COVID recovery | McKinsey. (n.d.). Retrieved May 1, 2022, from https://www.mckinsey.com/business-functions/operations/our-insights/how-retail-can-adapt-supply-chains-to-win-in-the-next-normal

Best practices to optimize the retail supply chain | McKinsey. (n.d.). Retrieved May 1, 2022, from https://www.mckinsey.com/industries/retail/our-insights/ten-steps-retailers-can-take-to-shock-proof-their-supply-chains

Understanding Retail Consolidation | C.H. Robinson. (n.d.). Retrieved March 4, 2022, from https://www.chrobinson.com/en-us/resources/retailconsolidation/

Zhang, G. P. (2003). Time series forecasting using a hybrid ARIMA and neural network model. *Neurocomputing*, *50*, 159–175. https://doi.org/10.1016/S0925-2312(01)00702-0