Rationalizing Inventory: A Multi-Echelon Strategy for Safety Stock Justification

by

Yu-Ta Chen
BBA. International Business, National ChengChi University

and

Rohan Alexander
BSc. Mechanical Engineering, University of Michigan - Ann Arbor

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 2023

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

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

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

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

Rationalizing Inventory: A Multi-Echelon Strategy for Safety Stock Justification

by

Yuta Chen

and

Rohan Alexander

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

ABSTRACT

This work presents a multi-echelon inventory optimization model for a manufacturing company to evaluate optimal inventory levels for a selection of products and their sub-components. The guaranteed service model is employed to identify potential improvements in inventory allocation while maintaining service levels. The model's results are compared with the company's current inventory policies to provide insights into the effectiveness of the proposed approach. First, demand forecasting was conducted for the selected products, and the results showed a close match with the company's existing forecasts. Next, a multi-echelon inventory optimization model was formulated using bill of materials, component lead times and standard costs. The model was optimized in Python to minimize total inventory holding cost, with constraints on service level, service time, and bounded demand. The model's output suggested an "all-or-nothing" type of inventory policy, wherein a stage either maintains zero safety stock or holds the maximum permissible safety stock. The results of the model revealed that 54% of the analyzed sub-components do not require any safety stock to be held. Additionally, the model proposes pooling inventory in stage 0, which is the finished product stage. The true financial impact of the model's results is difficult to gauge, because only a small portion of the product portfolio was used in the optimization. Potential areas for future work include investigating the impact of phantom stock removal, applying the stochastic service model to this problem, and understanding the impact of multi-echelon modeling on supply chain resilience. The insights provided in this work can serve as a starting point for manufacturing companies aiming to optimize their inventory policies and better manage their supply chains.

Capstone Advisor: Dr. Ilya Jackson
Title: Postdoctoral Associate

ACKNOWLEDGEMENTS

We would like to thank our sponsors for their flexibility and collaboration on this project. We would also like to thank our advisor, Dr. Ilya Jackson, for his constant support and motivation throughout the year. We would like to extend our deepest gratitude to Professor Sean Williams, who helped provide clarity on our model formulation. Finally, we would like to thank Professor Larry Snyder, who helped guide us through the technical challenges of implementing our model.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

1. INTRODUCTION

Inventory levels are critical to businesses trying to stay competitive in markets where affordable and timely order fulfillment is a basic expectation. Inventory shortages during the COVID pandemic resulted in long lead times and low fill rates for many businesses (Naughton, 2020). In contrast, a surge in inventory due to the bullwhip effect caused by the pandemic has forced many retailers to offer deep discounts to offload excess goods and reduce warehouse expenses (Nassauer & Terlep, 2022). Clearly, it is challenging for companies to balance between customer service levels and inventory costs. Maintaining appropriate inventory levels is important because it serves as a buffer against uncertainty and disruptions - but this comes at a hefty price of increased holding costs and working capital. From a customer perspective, storing large amounts of finished product inventory at all times seems like the obvious choice. But the solution is not that simple, especially for companies that manage thousands of varied products and an even larger plethora of sub-components needed to manufacture those products. The stock of each component must be accounted for: a lack of any one of them could result in downstream delays that ultimately propagate to the customer.

1.1 MOTIVATION

Our project sponsor, Water Company Inc., manufactures and sells water measurement devices for industrial, process, and laboratory use (Water Company Website, 2022). Their sites produce and store a variety of water instrumentation products such as handheld thermometers and drinking water analyzers. Water Company is in the process of revamping its supply chain strategy to better meet customer needs. Currently, the company does not employ a quantitative model to justify the amount and selection of finished products and sub-components it stores on site; instead, it often relies on heuristics or enterprise resource planning systems to make this decision. The result of this strategy may be higher than necessary inventory levels for certain products. High inventory levels occupy working capital and limit cash flow, causing poor return on assets.

Past State of S&OP

Water Company exclusively used demand forecasts to push inventory through their supply chain. No other product or component feature was considered when making inventory decisions. For a business of this scale, such a strategy may result in sub-optimal inventory levels, which makes managing inventory much more challenging.

Current and Future State of S&OP

Water Company uses a cloud-based forecasting system to predict demand for its SKUs. If a particular SKU has highly volatile demand forecasts, managers manually remove any occurrences of extreme demand from the prediction. Once outliers are removed, the demand forecast data is transferred to an ERP program which calculates base stock levels for finished product and raw material, and makes production order decisions automatically based on customer demand. Additionally, SKUs are now classified based on certain features (discussed in section 4), allowing the company to have custom order fulfillment strategies for different products. While custom fulfillment strategies add complexity to supply chain decisions, the company's ERP system is able to account for this when calculating inventory levels and placing production requests. However, the ERP system does not provide a quantifiable guideline for inventory levels of raw, finished and semi-finished goods: it is somewhat of a 'black box' system for supply chain planners. Therefore, the Water Company would like to determine whether the output from the program is optimal based on their cost and service level needs, and whether there is room to incorporate other variables that they have not considered yet.

1.2 PROBLEM STATEMENT AND RESEARCH QUESTIONS

As discussed in section 1.1, the company uses product segmentation to manage its inventory. SKU's are segmented based on economic impact and volatility, and corresponding inventory levels are assessed by managers at a frequency based on their criticality to the business. With thousands of complex SKUs, varied lead times and fluctuating demand to manage, Water Company's key problem is determining optimal inventory levels at their sites to satisfy customer orders in a timely manner, while minimizing costs.

Based on the business problem, our research question is "How can Water Company determine optimal inventory levels of raw, semi-finished, and finished inventory using a robust demand forecast?" To answer this question, we need to determine:

1) What factors are critical in determining inventory holdings for Water Company's SKUs?
2) How can Water Company use these dynamic factors to quantitatively determine inventory levels for its SKUs?
3) What is the impact of the new inventory policy on inventory cost and service level?

1.3 SCOPE: PROJECT GOALS AND EXPECTED OUTCOMES

The project's overall goal is to provide Water Company with an assessment of their current inventory levels, using a quantifiable model. This model will be applicable to other sites and product lines outside of this case study.

The success criteria for our project goals include:

1) Determine whether current demand forecast accuracy has room for improvement, recommend a new forecasting approach if true
2) Identify appropriate inventory modeling approach to handle raw, semi and finished goods

3) Quantify and optimize appropriate inventory levels using selected model and robust demand forecast

Prudent inventory management relies on accurate demand forecasts, given that appropriate inventory policies are built upon these forecasts. Though a forecast is always incorrect to some degree, we can work on minimizing forecast bias to set a baseline to start with. One hypothesis is that the Water Company does not have a reliable demand forecast, especially for highly volatile SKUs. As a result, their corresponding inventory policy cannot generate a workable guideline for inventory levels. After verifying the demand forecast accuracy, we can optimize inventory levels for their product, sub-components and raw materials. Alternatively, it is possible that the company's existing inventory policy accounts only for finished products and does not optimize for cost at the sub-component level. Another hypothesis is that the company treats each component as a single stage of inventory, thereby ignoring the implications of lead time and cost of components in neighboring stages.

Our findings indicate that the alternative hypothesis was correct. The ERP system calculates safety stock for each component independently, which is inherently sub-optimal for total system cost. Additionally, while the company has documented inventory policies and segments even at the component level, safety stock levels are sometimes determined by heuristics. For example, one segment of products has a safety stock target of yearly demand divided by 12, which is an overly simplistic way to determine inventory levels since it ignores the impact of lead time and cost. The project scope covers a small segment of 29 SKUs.

Expected outcomes of the project include:

1) A replicable demand forecast model
2) An inventory optimization model that provides optimal inventory levels for raw, semi and finished products

2. STATE OF THE ART

To determine the best method for optimizing Water Company's inventory levels, we focused our review of the literature on four core areas:

1. Demand Forecasting
2. Segmentation
3. Inventory Policy
4. Multi-Echelon Inventory Optimization (MEIO)

2.1 Demand Forecasting

Demand forecasting is a critical process in supply chain management that involves predicting the future demand for a product or service based on historical data, market trends, and other relevant factors. Accurate demand forecasting is essential for ensuring that a company maintains optimal inventory levels, minimizes waste, and meets customer demand in a timely and efficient manner.

2.1.1 Forecasting and its Impact on Setting an Inventory Strategy

To address Water Company's inventory issues, we must investigate their demand forecast first, because a robust forecast is the most important input for inventory policy. Under the inventory policy of order-to-level (S), S is decided by average demand and demand fluctuation. Also, many studies have revealed demand's impact on inventory investment and inventory cost. First, forecast accuracy has a significant impact on inventory investments (Bonney, 2009; Fisher, 2000). Second, low forecast accuracy is connected with overstock (Bonney, 2009). Third, whenever forecast error decreases, inventory cost always reduces (Jully Jeunet, 2005).

2.1.2 Model Selection

The first key issue in demand forecasting is which forecast method to use. A proper forecast method is important given that demand misspecification can lead to higher inventory cost (Badinelli 1990). To find the best prediction models, we need to identify demand patterns for each SKU (Erjiang, 2022). We refer to Erijian's model selection framework. First, we divide demand patterns into four categories: smooth, intermittent, erratic, and lumpy according to the squared coefficient of variation (CV^2) and the average inter-demand interval (ADI). Second, we build a pool of models. Considering that Water Company's project scope only includes SKU with relatively stable demand, time series and regression models might be good starting points. Detailed models might include naïve, seasonal naïve, single exponential smoothing (SES), Holt's linear exponential smoothing, damped trend exponential smoothing, Theta, 4-Theta, and simple combination of univariate models. Third, we use the dynamics weighting strategies in the literature review to predict. Lastly, we review the effectiveness of each prediction model with Water Company's historical data.

2.1.3 Evaluating Forecast Accuracy

The second key issue is how to minimize forecast error. We are interested in forecast accuracy not only because it is important to inventory cost as we mentioned above but also because historical performance can generate the description of future demand to make decision for resources allocation and help us improving the process (Silver et al., 2016)
There are several ways to evaluate forecast accuracy, such as mean squared error (MSE), mean absolute deviation (MAD), and mean absolute percent error (MAPE). For our use case, MAD is less important given that there is no need for computational simplicity due to the limited number of SKUs. The benefit of MAPE is that it is a percentage and therefore not affected by demand magnitude, but MAPE is not suitable when demand is very small (Silver et al., 2016). In Water Company's case, either MSE or MAPE may serve as a good matrix of forecast accuracy.

2.2 Product Segmentation

Segmentation is a method used by companies to classify their products based on a business metric of their choosing. Segmenting products allows companies to devote more time to goods that are critical to the business. It also allows managers to create policies specific to each segment, without delving into details for each product. This ‘divide and conquer’ method is especially useful when companies manage thousands of commodities, making it time consuming to treat each product with the same level of diligence. (Lopez et al., 2013) Segmentation is based on the Pareto Principle, which states that roughly 80% of consequences are generated by roughly 20% of causes. Applied to the business domain, the Pareto Principle implies that most of the revenue generated by an operation comes from a small fraction of its products (Rungtusanatham et al., 2010, #). It is only logical that managers treat their most valuable few products differently than the rest.

2.2.1 ABC Segmentation

ABC is one of the most widely used segmentation strategies. It sorts products into three categories (ABC) based on their value to the business. Historical sales data is aggregated and sorted in descending order to obtain the revenue contribution of each SKU (Hoffmann et al., 1991). This process can be performed during a fixed or dynamic review period depending on the business need. For example, a fashion company may want to redefine its segments at a higher frequency because of the seasonality of its products - a shirt that was popular last year may be out of fashion this year. It is important that only demand during the review period must be considered during segmentation, since demand can be dynamic, and SKUs can move across segments as a result.

Next, SKUs that account for 70%-80% of revenue are grouped into segment ‘A’. This segment typically contains 15-20% of the total number of SKUs. Following the same process, segment ‘B’ usually contains 30-40% of SKUs that account for 15-20% of revenue, and segment ‘C’ contains the rest.

2.2.2 ABC/XYZ Segmentation

Water Company uses ABC/XYZ segmentation to classify their products. This grouping is simply a combination of ABC coupled with a similar classification based on demand volatility, resulting in nine different product segments. For example, segment AZ contains products that present the most financial value to the business, but also have volatile demand. Segment CX, on the other hand, represents products that do not generate much value for the business, but have stable demand. Figure 1 shows an example of the ABC/XYZ classification technique. Segmenting products using an ABC/XYZ system allows managers to get more granular with their decision making. In the supply chain domain, variables such as safety stock, order quantity and frequency can be set based on the segment a product belongs to.

FIGURE 1:

An example of ABC-XYZ segment thresholds (SAP, n.d.)

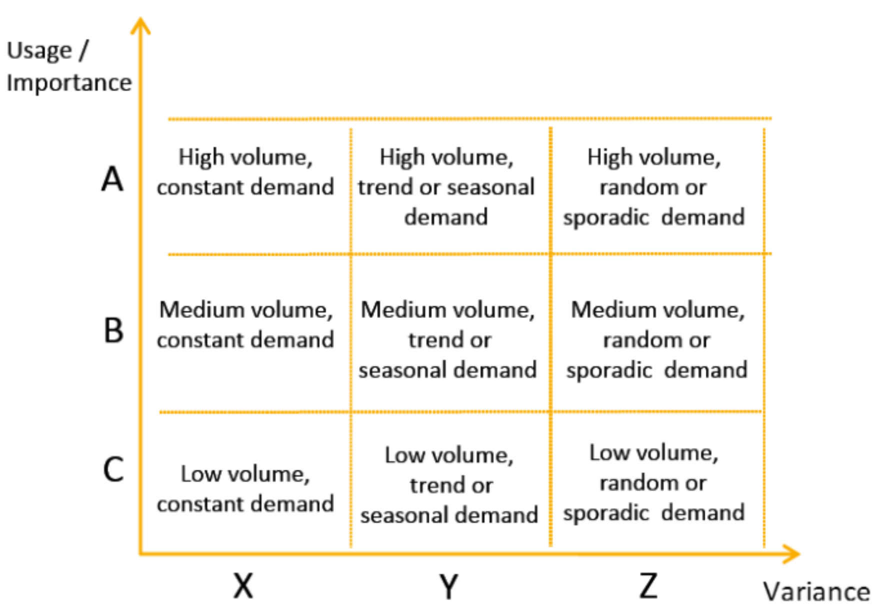

2.3 Inventory Policy

Inventory is critical to companies trying to protect against uncertainty in demand and supply of its products. Excess inventory could result in high holding costs for the business, while inadequate inventory levels could result in unfilled customer orders and lost revenue. Companies therefore employ inventory management systems to reduce inventory costs while satisfying customer orders. An inventory policy comprises two essential elements: how much to order and when to order. The inputs that go into determining this policy can be as simple as average demand - or could extend to additional factors like lead time and service level. Below we review some commonly used inventory management policies.

2.3.1 Economic Order Quantity (EOQ)

EOQ is one of the most elementary inventory policies used to determine what the optimal order quantity for a product should be. Here, optimality is based on minimizing inventory holding and ordering costs (Harris, 1913). The formula for EOQ is the following:

$$EOQ = \sqrt{\frac{2kD}{h}}$$

where k = ordering cost, D = demand and h = holding cost

The EOQ model is simplistic because it assumes constant demand and lead time - both of which are unlikely to occur in a real operation.

2.3.2 Periodic Review Model (R,S)

Periodic review policies are ideal for businesses that order goods in regular intervals. In this case, the policy determines 'S' - which is the order up to level, while R is the review period already set by the company. The formula to calculate 'S' is as follows:

$$S = \mu_{DL+R} + k\sigma_{DL+R}$$

where μ_{DL+R} is the expected demand over the lead time + review period, k = safety factor based on desired service level and σ_{DL+R} = standard deviation of demand over lead time + review period (Silver et al., 2016).

Increasing the lead time or review period will increase the average inventory costs for the company. This is because the average amount of inventory held will increase to account for longer wait times between replenishment, thus incurring additional costs to the business.

2.3.3 Continuous Review Model (s,Q)

Continuous review policies are ideal for businesses with systems that can constantly monitor inventory levels. In this model, an order of quantity 'Q' is placed every time inventory levels drop below a level 's'. Typically, 'Q' is calculated using the EOQ formula discussed above.

To calculate the reorder point s, the following formula is used (Silver et al., 2016):

$$s = \mu_{DL} + k\sigma_{DL}$$

This is very similar to the formula for order up to level in the periodic review model above. The only difference is review time = 0 since inventory levels are constantly monitored.

Similar to the periodic review policy, increasing lead time will increase the reorder point to account for the higher average and standard deviation of demand during lead time. This in turn will increase inventory costs. Increasing the safety factor 'k' will have the same effect by increasing the safety stock held by the company to achieve a higher service level. Thus, these policies pose a trade-off for managers between customer service levels and cost of safety stock held.

2.4 Multi-Echelon Inventory Optimization (MEIO)

One of Water Company's goals is to quantify and optimize appropriate inventory levels that balance raw materials, semi-finished materials, and finished goods. However, the inventory policies mentioned in section 2.3 have limitations: they can only address one product or component at a time. To fix this limitation, we bring in multi-echelon inventory optimization (MEIO).

MEIO is a supply chain management strategy that involves managing inventory across multiple levels or echelons of the supply chain. In a typical supply chain, inventory is held at various stages. These different stages can be raw materials, semi-finished materials, and finished goods within a BOM (bill-of-materials), or they could be the manufacturer, distributor, and retailer levels. MEIO aims to ensure that the right amount of safety stock is held at each echelon, to minimize inventory costs while ensuring that customer demand is met (Snyder et al., 2019).

Key questions to answer are which stages should hold safety stock, and how much to hold. It is not necessary for all stages to hold safety stock. The stages holding safety stock serve as buffers to absorb the demand uncertainty. Identifying which stages should hold safety stock is a strategic problem known as strategic safety stock placement problem (SSSPP), because it is expensive to change the design frequently.

MEIO shares similar assumptions with the Periodic Review Model (R,S): that time is infinite-horizon and each echelon follows a base-stock policy. Each stage quotes a committed service time (CST) to its downstream echelon within which it will deliver orders from the downstream echelon. Each stage is required to offer 100% service to downstream stages, which means that each stage must satisfy orders within CST no matter what. The concept of CST is the key to connect each stage and to decide safety stock. As the base stock increases, the stage can offer a shorter CST to the next stage. The inbound service time of a downstream stage (close to finished products) equals the outbound service time of an upstream stage (close to raw materials). The time needed to finish whatever action needed in a stage is the lead time or the process time of the stage. A visualization of these times can be seen in figure 2.

FIGURE 2

Depiction of service and process times of each stage

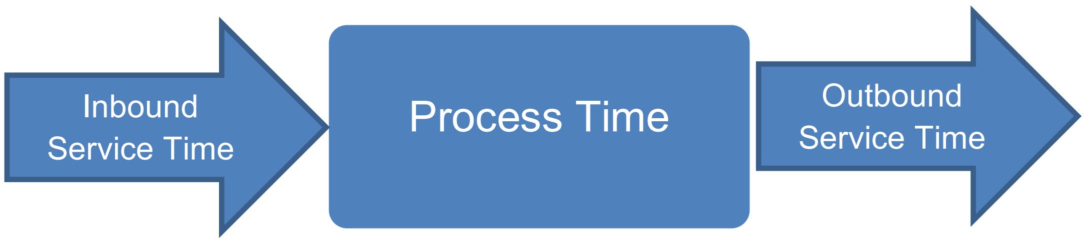

The decision variable of the MEIO is the service time because service time decides the inventory level of a stage. Whenever we decide the outbound service time of a stage, we have the inbound service time of the following downstream stage too. Inbound service time of stage i SI_i plus process time of stage i T_i minus outbound service time of stage i S_i is what we call net replenishment time. Net replenishment time is one of the factors deciding the inventory level. In the MEIO model, safety stock at stage i is proportional to the net replenishment time of that stage. In an MEIO model, safety stock at stage i equals to $h_i\, z_\alpha \sigma_i \sqrt{SI_i + T_i - S_i}$,where $z_\alpha \sigma_i$ is z standard deviations above the mean for some constant α, σ_i equals to standard deviation of demand in stage i, and h_i is the unit holding cost.

In reality, it is challenging for companies to deal with highly volatile demand. Therefore, some MEIO models assume that demand is bounded in any time interval, meaning that demand fluctuates within certain ranges. We model the demand by assuming the demand is normally distributed and truncating the right tail of the demand, meaning that we ignore any demand greater than certain standard deviations above the mean, which is the z_α we mentioned above. For realized demand above the bound, one possibility is that the company will handle them by other approaches such as outsourcing and overtime shifts.

In terms of stage definition, a node will be considered a stage only if it can hold inventory. If one stage processes the materials and then passes to the next stage, such a stage should be combined with another stage that can hold inventory and be viewed as just one stage.

Within MEIO, there are two primary models: guaranteed-service model and stochastic-service model. The result of either model is a set of base stock levels, but the two models require different inputs to determine the base stock level (Snyder et al., 2019).

Kimball introduced the guaranteed-service assumption in 1955, which was later reprinted as Kimball (1988). Simpson (1958) applied this assumption to serial systems, while Graves (1988) discussed the optimization of safety stock resulting from this assumption. Inderfurth (1991), Minner (1997), and Inderfurth and Minner (1998) developed dynamic programming (DP) approaches for distribution and assembly systems. Graves and Willems (2000) expanded on this by applying DP to tree systems, and Magnanti et al. (2006) and Humair and Willems (2011) extended DP to general networks that include undirected cycles.

The structure of multi-echelon networks can vary, and their topology significantly impacts the analysis and optimization of the system. Topologies include serial systems, assembly systems, distribution systems, tree systems, and general systems. In serial systems, each echelon contains exactly one stage. Assembly systems mean that each stage has at most one successor. Tree systems combine features of both assembly and distribution systems, with each stage potentially having multiple predecessors and successors. However, tree systems are characterized by the absence of any undirected cycles. In Water Company's case, even though the network of BOM (bill-of-materials) is typically described as assembly systems, it can also be viewed as one type of spanning tree whenever there are multiple predecessors and successors at each stage. To enhance the practicality of our research, we will use the dynamic programming algorithm from Graves and Willems (2000) for tree systems to find the optimal safety stocks.

Graves and Willems (2000) proposed a model and algorithm for multi-echelon inventory optimization, which is known to run in pseudo polynomial time. However, in 2004, Lesnaia devised an algorithm with a polynomial-time complexity of O(N3), where N corresponds to the quantity of stages comprising the network. For more general systems that include undirected cycles, the problem is NP-hard (non-deterministic polynomial-time hardness), as indicated by Chu and Shen (2003) and Lesnaia (2004). Magnanti et al. (2006) presented a solution method based on integer programming techniques to address multi-echelon inventory optimization problems in general systems. Humair and Willems (2011) developed exact and heuristic algorithms that extend the DP algorithm introduced by Graves and Willems (2000) to general systems. Additionally, Humair et al. (2013) extended the approach to account for stochastic lead times. Furthermore, Graves and Schoenmeyr (2016) studied the effects of capacity constraints in the context of multi-echelon inventory optimization.

To summarize, our literature review covered demand forecasting, product segmentation, inventory policy, and inventory optimization. The most critical aspect of this review is multi-echelon inventory optimization, since this will ultimately provide a quantifiable justification for inventory levels. Within MEIO, the guaranteed service model approach is most applicable to our sponsor's objective, because of their service level requirements. This choice of model will be explained in further detail in section 3.

3. METHODOLOGY

The goal of this capstone is to provide the sponsor company with a quantifiable justification of inventory levels for raw, semi and finished products. In section 2, we discussed several viable methodologies to calculate base stock levels. Of these, we believe multi-echelon inventory optimization is most appropriate for the company's problem, because the bill of materials provided to us closely resembles a tree system. Specifically, we will use the guaranteed service model to optimize safety stock levels at each raw, semi and finished product stage. The methodology for this capstone will be divided into four stages, as depicted in figure 3 below.

FIGURE 3

Flowchart of methodology

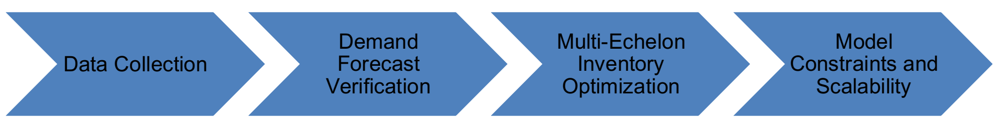

3.1 Data Collection

To determine the optimal inventory levels for Water Company, we started by focusing on a small sample of their SKUs. The initial selection of SKUs has relatively stable demand patterns to allow us to build and test the model before scaling to SKUs with sporadic demand. To begin our assessment, we collected the following data:

1. Real and forecasted monthly demand for selected SKUs starting from 2019

2. Company’s service level requirements for its products
3. Process time for production and delivery of raw/semi-finished/finished goods
4. Bill of Materials for selected SKUs
5. Historical inventory levels of selected SKUs

3.2 Demand Forecast Verification

We verified the accuracy of the company’s demand forecasts to ensure that their current forecasting software is performing optimally. We verified forecast accuracy by comparing the root mean square error (RMSE) of their forecasts with the RMSE of forecasts we generated using the modeling method chosen by their software. The forecasting software used by the company generates a new monthly forecast for each finished good at the beginning of every month. While the software displays which model was used to create predictions, the parameters used are auto generated for each forecast. To determine whether the accuracy of these predictions for each SKU could be improved, we trained the same models chosen by the software on 24 months of historical monthly demand data, and tested it on 12 months of predictions from September 2021 to September 2022. All the models chosen by the software were statistical. Model parameters were tuned to achieve lowest possible error - although these parameter values may not have matched those chosen by the software. In addition, we used a ‘Prophet’ model with auto selected parameters as an alternative forecasting approach to the models selected by the software. While we could have used several other forecasting models, we chose to shift our focus to inventory optimization, since this is what will ultimately allow us to validate the company’s inventory levels.

3.3 Multi Echelon Inventory Optimization

Once the demand forecast was verified, we developed an inventory optimization model to address the company's challenge of justifying inventory levels at the raw, semi-finished and finished stage. A multistage model is a more accurate representation of the company's value chain, since it accounts for intermediary production stages as opposed to a single stage model which considers only the finished product stage. We decided to proceed with the guaranteed service model, since it is commonly implemented in practice (Eruguz et al., 2016).

3.3.1 Model Selection

Within the scope of multi-echelon inventory optimization, there are two widely used approaches: guaranteed service and stochastic service. The stochastic service model (SSM) accounts for stock outs at each stage by applying a penalty for each time period there is no inventory available. In contrast, the guaranteed service model (GSM) assumes no stock outs can take place, since demand is bounded. While the guaranteed service approach employs less realistic assumptions because of this bound, it is easier to implement and closely matches our sponsor company's needs, since it guarantees that the inventory positioning at each stage is enough to meet customer service level targets. Additionally, any realized demand over the demand bound can be met using other strategies, such as increasing production capacity. Further, Klosterhalfen and Minner (2010) found that the GSM base-stock policy is more cost-effective than the SSM policy for systems with long lead times service levels - although this research was more specifically applied to warehouse systems. For the reasons stated above, we decided to proceed with using the guaranteed service approach to tackle this multi-echelon inventory problem.

3.3.2 Model Formulation

To set up the model, we performed the following four steps:

1. *Removed irrelevant components in the BOM (bill of materials):*

The BOM includes irrelevant components including phantom components, irregular components, and components with zero accounting cost. Phantom components are virtual stages that do not physically exist or exist only during the assembly phase. These components are still created virtually to serve other purposes such as accounting. One concern about removing phantom stages is that they have zero stage time, yet a direct cost was added. That added cost is usually the overhead cost spread to the assembly. In our case, materials at levels closer to finished products are not phantom, so eventually these higher-level stages include all values added in lower level. Therefore, such a concern is not serious in our case.

Irregular components such as water are not tracked in the inventory level. The MEIO model requires the cost of goods sold or values added to each stage as an input to minimize inventory holding cost, so if a component has zero cost, the model will naturally stock inventories at this stage. However, such a result does not reflect the true inventory holding cost.

2. *Structured the BOM*

Many materials such as screws are used multiple times at different levels of the BOM. We treated them as different materials and aggregated the result at the end. The reason we took this approach is because the topology of the BOM might have a circle. One key definition of the spinning tree system is N stages with N-1 arcs. If the topology contains circles, the algorithm for the tree system might not apply because the topology is no longer N stages with N-1 arcs. Figure 4 provides a visual representation of what one such BOM structure looks like.

FIGURE 4

BOM structure of one of the SKUs under scope, created using networkX

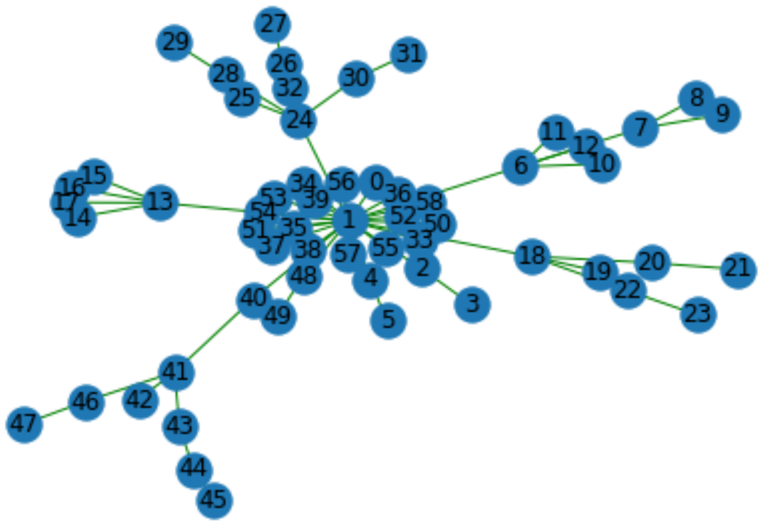

3. *Re-labeled stages*

The DP (dynamic programming) algorithm requires stages following a labeling rule that each stage (other than stage N, where N is the number of stages) has exactly one adjacent stage with a higher index. In the topology of a BOM, it is fairly easy to label materials compared to other tree systems. All we must do is label materials of higher steps with smaller indexes. If we look at the network displayed in Figure 5(a) and apply the process, we obtain the reorganized network shown in Figure 5(b). It is worth noting that in this network, except for stage 6, every stage has a single neighbor with a higher index, whether upstream or downstream.

FIGURE 5

Relabeling the network

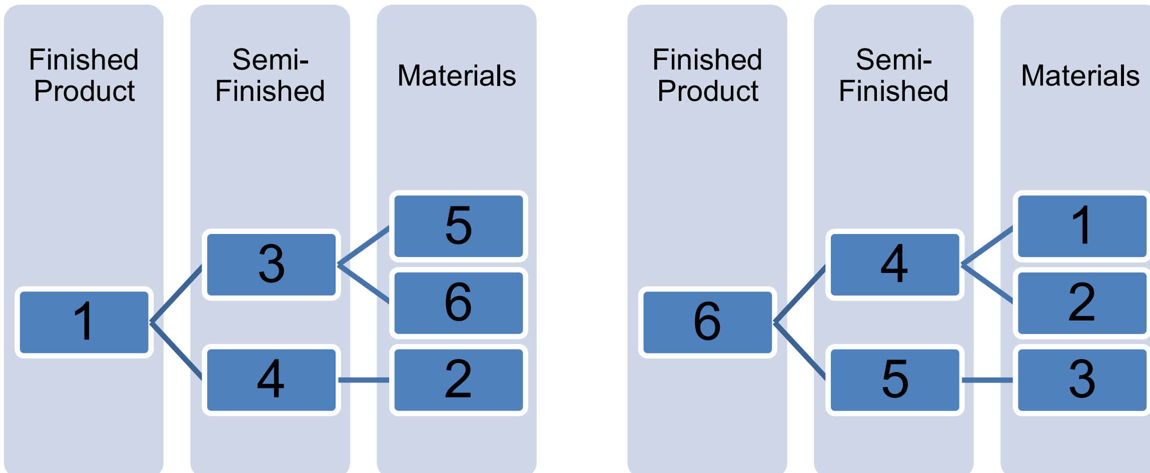

(a) Original network. (b) Relabeled network.

4. *Input service time and review period based on segmentation*

Outbound service time for each finished product was then input into the model. This was unique to each finished product, based on the segment it belonged to. Similarly, the purchasing review period - also determined by segmentation, was added as a model input. This was added to the process time of each stage. Assumptions about these inputs will be discussed in section 3.3.2.

5. *Formulated objective function as total holding cost across all stages*

What we want to know is how much inventory each stage should hold without changing the design of the network. Given that the network is unchanged, the cycle stock will not change with different configurations of the network, where ‘configuration’ means different safety stock allocation. Safety stock is what really changes under different scenarios. Therefore, our objective function is minimizing total inventory holding cost at every stage:

$$\sum_{i=1}^{N} h_i \, z_\alpha \sigma_i \sqrt{SI_i + T_i - S_i}$$

6. *Formulated constraints for process time*

The first constraint is that net replenishment time, $\sqrt{SI_i + T_i - S_i}$, must be greater or equal to zero. The second constraint is that whenever a stage has multiple upstream stages, the inbound service time of this stage must be the greatest value among outbound service time of all upstream stages. Except for the above two constraints, there are some given variables. First, the outbound service time of stage N is requested by customers, so it depends on what Water Company's customers request. Second, the inbound service time of the most upstream stages, which are the materials, will be combined with the process time, meaning that we assume the lead time needed to obtain the materials is the process time of the stage and inbound service time equals zero.

3.3.2 Model Assumptions

Several key assumptions were made when formulating the raw data to fit our model.

1. *Manufacturing Process Time:* Assembling process time is 1 day for each level. This assumption was made because the Python package requires the process time to be an integer. We understand that we might slightly overestimate the inventory level because actual manufacturing process time is less than one day for most components, but we think this overestimation is manageable because the gap is small, and it is better to have a conservative estimate. It should be noted however, that while the actual process time of production is approximately 1 day, company records maintain a 3-day process time for most assembled components, because most components are made in large batches to take advantage of economies of scale.
   Also, inventory review cycle is added to process time to cover the demand variety during the between the review periods.
2. *Substitution of Standard Deviation:* We used RMSE of forecast error instead of demand standard deviation. The purpose of safety stock is to mitigate forecast error, which is

why the size of safety stock should be proportional to forecast error. We picked the monthly demand forecast in 2022 and converted the monthly RMSE into daily RMSE using the following formula:

$$Daily\ RMSE\ =\ Monthly\ RMSE\ \div\ \sqrt{\frac{365}{12}}$$

3. *RMSE of subcomponents:* One piece of finished product might require multiple pieces of materials, so we needed to incorporate the ratio of finished goods to materials into the RMSE of each material into the model. Finished product RMSE was converted to the sub-component level using the formula below:

$$Component\ RMSE\ =\ Finished\ Product\ RMSE\ \times\ Usage\ Ratio$$

   As an example, if 0.04 units of material A is required to produce 1 unit of finished product with forecast RMSE of 100, then the RMSE of material A is 100*0.04 = 4.

4. *Purchasing Review Cycle:* The purchasing team reviews inventory and places orders for all purchased materials on a frequency based on the segment of the component. Each segment has a range of yearly review cycles, and the true frequency within this range is decided by purchasing managers. For our modeling, we chose to use the upper bound of this review range. For example, purchase orders for products in the AX segment are placed 12 - 24 times each year. We assumed that this frequency is 24 months, which would mean our review cycle = 365/24 = 15 days. Additionally, due to modeling constraints, we assumed that the review cycle of each component is the same as its parent finished product.

5. *External Outbound Service Time:* Like the review cycle, service time to the customer is based on the segment a finished product belongs to. For example, if the product is in the AX segment, it has an external customer service time of 1-2 days. In this case, we chose to use the lower bound of these time ranges to provide a best-case estimate.

6. *Capacity Constraints:* For the purpose of this model, any capacity constraints faced by the company were ignored. An important assumption of the guaranteed service model is bounded external demand. If demand crosses this threshold, we assume that the company can meet this demand by expediting orders or increasing throughput. The same goes for storage capacity: if the model proposes an inventory increase, then we assume that the company has enough storage space to accommodate for this.

3.3.3 Model Optimization

After the model was formulated, it was optimized in Python to minimize total inventory holding cost. The resulting stage process times were then used to calculate optimal safety stock levels for each component in the BOM. We used an open-source python package called ‘Stockpyl’ which embedded the algorithm introduced in Graves and Willems (2000) to assist with the optimization. The dynamic program iterates through every combination of potential outbound service time from the lower index to find the optimal solution that minimizes total system holding cost.

4. RESULTS

In this section, we will present the results of our demand forecasting and inventory optimization models. The results will be benchmarked against the company's existing demand forecasts and inventory levels respectively.

4.1 Demand Forecasting

The error of our forecasting models conducted across all 29 selected SKUs was on average within 3% of the company's forecast error. The relative model performance is depicted in Figure 6.

With such closely matching forecasts, we do not believe there is much room for improvement in the company's current modeling approach. However, this is an important initial verification, because an error-prone forecast could severely impact future inventory decisions. Substantiating the demand forecast allowed us to shift our focus to the next stage of our methodology: inventory modeling.

FIGURE 6

Comparison of our forecast and the company's forecast

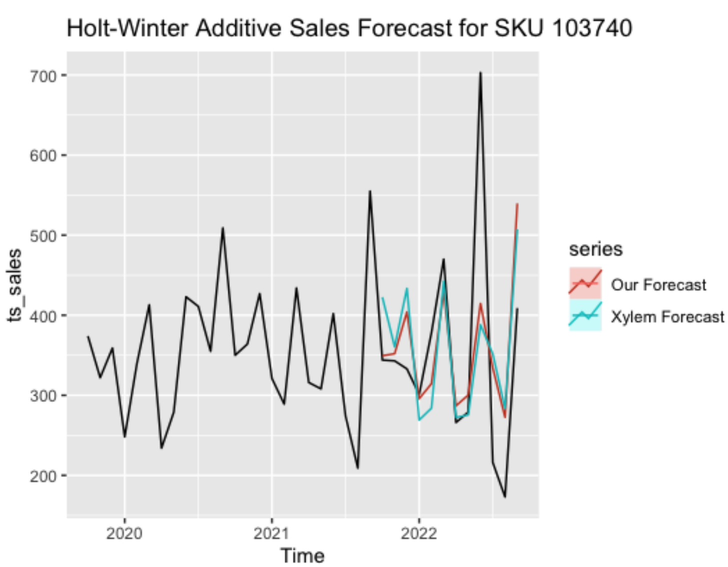

4.2 Multi Echelon Inventory Optimization

The results of our model were committed service times for each stage. These were converted to proportional safety stock and cycle stock levels of each component, which were then used to estimate average inventory levels and benchmarked against the average inventory levels for that SKU over the past year.

Out of the 522 subcomponents analyzed, the model estimates that 54% of these components do not require any inventory to be held at all. Additionally, the model proposes pooling inventory in stage 0. A more detailed analysis of these results will be presented in the sections below.

4.2.1 All or Nothing Result

Given the nature of the algorithm, the inventory policy employed at each stage is of the "all-or-nothing" type, wherein the stage either maintains zero safety stock and quotes the maximum achievable cycle service time (CST) or holds the maximum permissible safety stock and quotes zero CST.

From our result, we can see that there are 285 materials holding zero safety stock among 522 materials and finished products, and the other stages serve as buffer stages to carry safety stock on behalf of the supply chain. Currently, the company only has 72 components that do not hold safety stock. Our solution will allow the company to reduce the complexity of managing safety stock across multiple SKUs.

4.2.2 Inventory Positioning

To interpret whether our model was pushing inventory upstream or downstream within the system's internal value chain, we calculated the total value of inventory in all component levels. Component levels numbered from 0 to 18, with 0 being the finished product, which is furthest downstream, and 18 being raw materials, which are furthest upstream. We then calculated the percentage of total system value each level held. From our result, we found that the model pushes inventory downstream (Figure 7 and Figure 8): specifically, it consolidates inventory at level 0 (finished product). Looking at table 1, we can see that level 0 has a high net replenishment time of 52 days, which dictates how much safety stock is held at that stage. This is because the committed service time to customer is low for the segment of products under scope. To achieve these service time constraints, along with the 90% service level constraint while minimizing total system cost, the company must pool inventory at this stage.

TABLE 1

Proposed Allocation of Inventory Value in Each Level

Level	*% of total safety stock value*	Level	*Average Net Replenishment Time (day)*
0	67.97%	0	52
1	2.83%	1	38
2	8.13%	2	52
3	1.77%	3	19
4	14.38%	4	18
5	0.87%	5	11
6	2.65%	6	32
7	0.05%	7	1
8	0.02%	8	17
9	0.05%	9	8
10	0.33%	10	17
11	0.10%	11	10
12	0.17%	12	18
13	0.10%	13	7
14	0.01%	14	17
15	0.40%	15	33
16	0.02%	16	29
17	0.05%	17	40
18	0.10%	18	56
Grand Total	100%	Average	11

FIGURE 7

Safety Stock Allocation Across Echelons

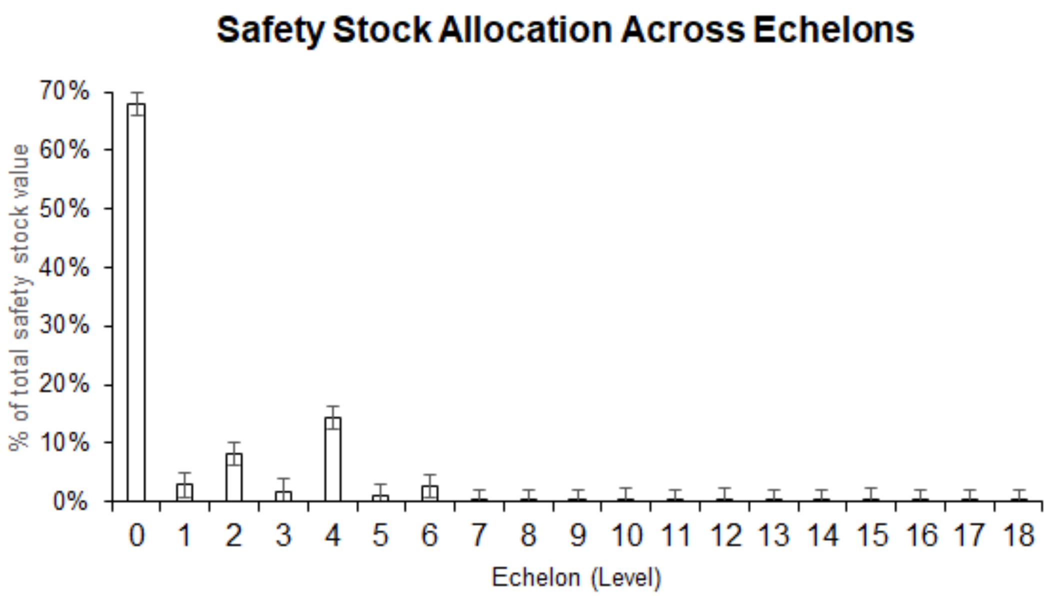

FIGURE 8

Net Replenishment Time Allocation Across Echelons

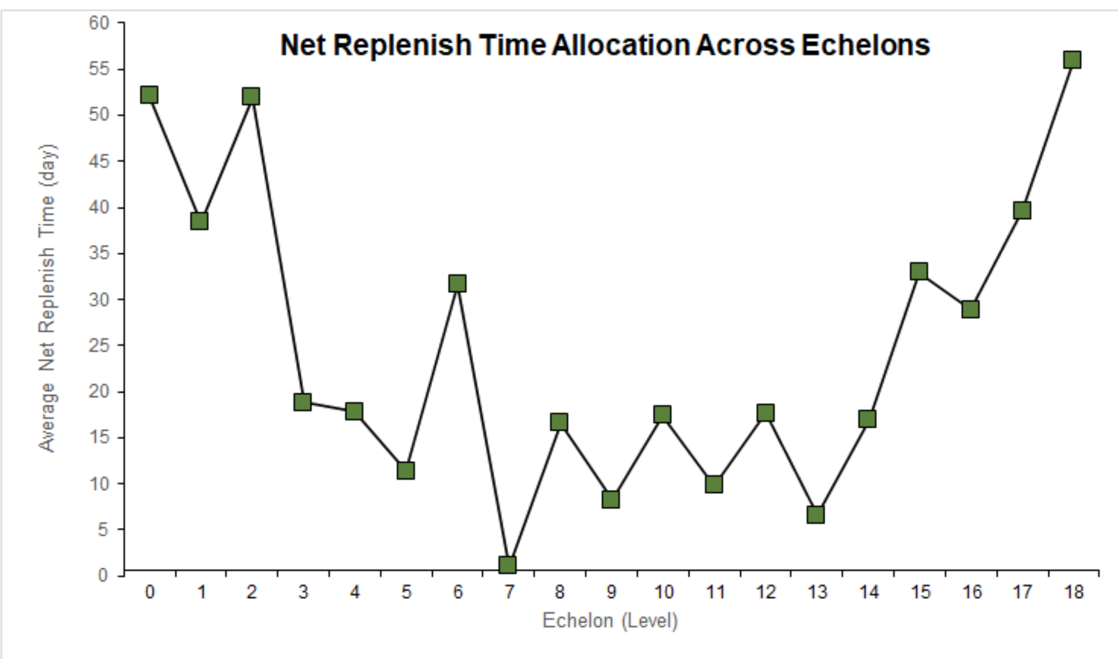

TABLE 2

Average Cost and Process Time for Each Sub-Component Level (process time includes review period)

Level	*Standard Cost (USD)*	*Process Time (Days)*
0	224	23
1	141	27
2	14	59
3	5	58
4	5	53
5	10	48
6	7	71
7	2	43
8	61	68
9	43	47
10	34	51
11	51	45
12	39	51
13	32	40
14	32	44
15	2	35
16	8	45
17	39	49
18	263	56

5. DISCUSSION

In this section, we will derive business insights from our results, and discuss how managers can use these insights to make better business decisions. Additionally, we will discuss the limitations of our model results and the implications of these limitations on the value chain of the product segment under scope.

5.1 Managerial Insights

This section will cover the business implications of model implementation. The model results will reveal which components should hold safety stock to minimize total inventory value while meeting customer service requirements. Managers can then use these results to decide whether the financial impact of implementing new base stock policies across its product lines aligns with the company's financial goals.

5.1.1 Sensitivity of Process Time and Review Period

Assuming a review cycle of seven days, table 3 illustrates the relationship between service time and inventory value. The second row indicates the inventory level as a percentage of the baseline, which is set to 0 days. While the review cycle may vary based on SKU segmentation, this table can serve as a reference point for Water Company in determining the optimal service time. For example, if the company wishes to move certain product families to segments with different service times, they can approximate the impact on working capital and holding cost using figure 9 and figure 10. Similarly, if the company wishes to adjust its service levels to meet changing customer expectations, table 4 can serve as an approximation of the financial impact. However, the model and sensitivity analysis must be performed on the entire selection of SKUs to accurately assess the impact of these changes.

TABLE 3

Service Time Sensitivity

Service Time(days)	0	7	14	21	28	35	42
Inventory value as % of inventory value with a service time of 0 days	100	89%	72%	62%	58%	56%	54%

TABLE 4

Service Level Sensitivity

CSL(%)	80	85	90	95	98	99
Safety stock value as % of safety stock value with a CSL of 80%	100%	123%	152%	195%	244%	276%
Average inventory value as % of average inventory value with a CSL of 80%	100%	109%	119%	135%	153%	165%

5.1.2 Inventory Pooling

Drawing from Section 4.2.2, "Inventory Positioning," it is evident that the proposed model recommends a consolidation of approximately 67.97% of the safety stock at level 0, corresponding to the finished product level. A comparative analysis between the model's projected inventory level and the actual inventory figures from 2022 in Table 5 suggests that the Water Company may need to augment its inventory levels for finished goods.

This projected increase can be attributed to two possible factors. Firstly, the segmentation of the chosen SKU necessitates a minimum review cycle of 15 days and offers a low service time to the end customers. This, in turn, could potentially prolong the net replenishment time. Despite upstream stages quoting a zero service time for finished products, the extended review period and the lower service time afforded to the end customers are immutable variables. In reality, existing finished product inventory levels may be too low to guarantee on time delivery to the customer.

The second contributory factor relates to the service level. The model assumes a service level of 90, whereas the actual service level could potentially be lower. Unfortunately, we lack historical data on service levels, making it challenging to validate this assumption. However, it would be intriguing to investigate whether the Water Company consistently maintains this projected service level.

TABLE 5

Comparison between Actual Inventory and Model Output

Level	*Actual Inventory in 2022*	*MEIO Model Result*	*Difference*	*Difference in %*
Value (USD)	$45,900	$130,058	+$84,159	+183%
Quantity (Units)	262	671	+409	+156%

5.1.3 Full-Scope Benchmarking

One of the benefits of multi-echelon modeling is that it considers the relationship between materials when deciding committed service times for each stage. For this reason, it is important for managers to only implement the calculated base stock levels for components that are unique to the SKUs used as inputs when the model was run. For example, if a component is used across 100 different finished products, all those products must be used as model inputs before making inventory decisions for that component. This can be a challenge for components that are standardized across all SKUs, since computation time will increase significantly. We recommend that the company run this model across all SKUs to get committed service times for each stage and use these CSTs to calculate base stock levels every time a new demand forecast is generated. The only time the model would need to be re-run is if any of the model inputs change. For example, new product introductions or lead time reductions would impact the model output. Here, managers can decide what frequency of model implementation makes most sense - based on the frequency of new product introductions, raw material alterations, cost savings projects, or any other changes that may significantly impact model inputs

5.1.4 No Operationalizing Needed

The output derived from our model is a series of base-stock levels. Notably, this occurs even though the decision variables within the guaranteed-service model are centered around cycle-service times (CSTs) as opposed to base-stock levels. Consequently, the applicability of the guaranteed-service model extends beyond scenarios where stages explicitly quote CSTs to each other, demonstrating its versatility in various operational contexts.

In short, Water Company can refer to the model result even when stages do not actually quote service time to one another in real operation. If the company chooses to implement model results, the only change they would be required to make is base stock levels in their ERP system.

FIGURE 9

Average Inventory Level Increase as a Function of Service Level (Baseline = 80%)

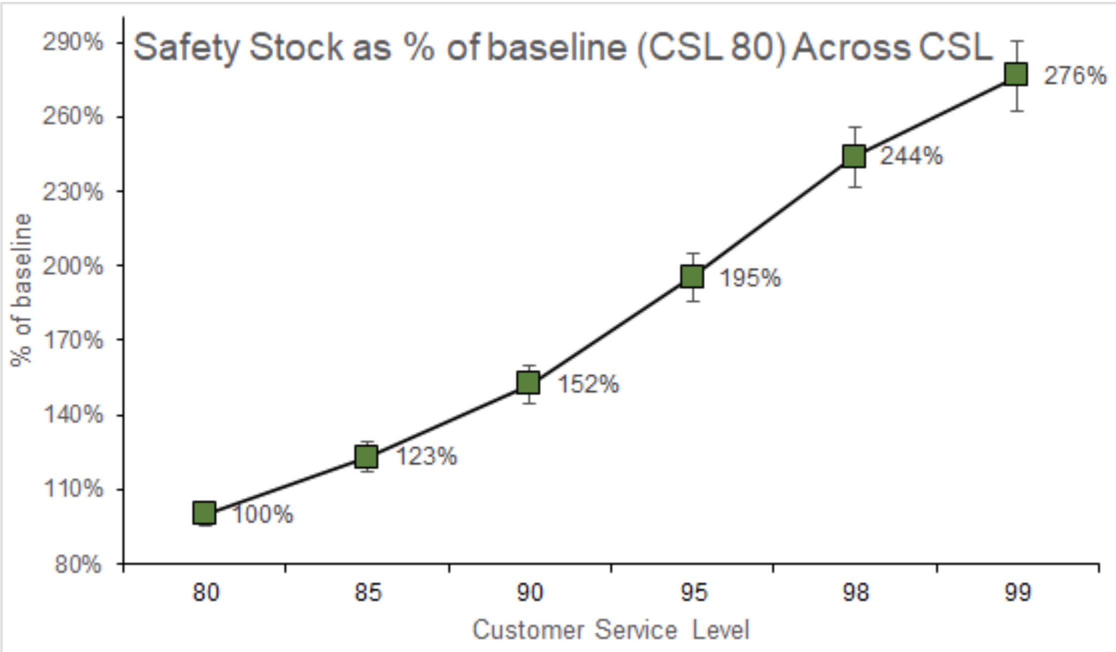

FIGURE 10

Average Safety Stock Level Increase as a Function of Service Level (Baseline = 80%)

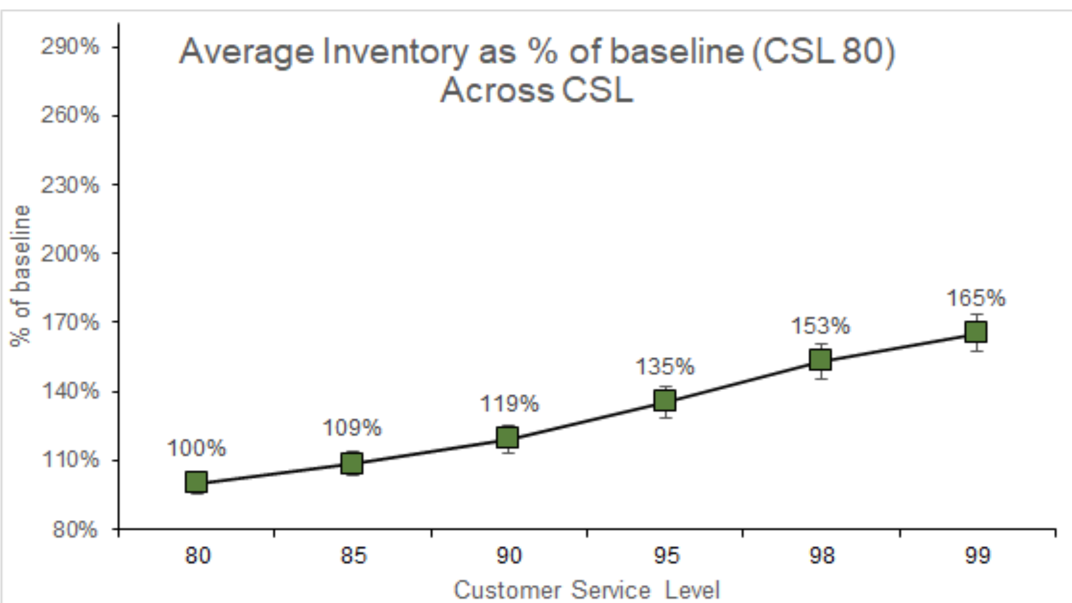

5.2 Limitations

The multi-echelon model utilized in our study exhibits several limitations that stem from underlying assumptions outlined in section 3.3.2, as well as its potential applicability within a business context. These limitations primarily arise because of the structure of BOM components and the model formulation itself, both of which are expounded upon within this section.

5.2.1 Syncing Service Times Across Materials

In many of the BOMs, the same material might be used in different semi-finished products. Whether we should treat the same material as different stages was a modeling challenge. If we treated them as the same stage and they all quoted the same outbound service time for different semi-finished products, then there is no problem. However, if they quoted different outbound service times for different semi-finished products, then it would be difficult to operationalize the result of the model, because the materials management team might not be able to set up different service times for the same materials. To fix this issue, we tried to sync the service time of the same material. We could either follow the largest or the smallest service time. If we synced with the largest service time among all service time quoted to next stages, it meant that we would be pushing safety stock downstream, because larger outbound service time reduces the net replenishment time and therefore reduces the safety stock of these particular materials. In contrast, if we synced with the smallest service time, we would be pushing safety stock upstream. In typical supply chain cases, stages keep adding value as we go downstream, so pushing safety stock upstream should create less inventory value and inventory holding cost. The issue with syncing service times is that the resulting changes in committed service times of neighboring nodes would not be captured in the model results. Consequently, the model would no longer guarantee optimality. Another possibility would be to pool duplicates into one stage. The drawback with this approach is that the BOM topology will no longer be a spanning tree, and our dynamic programming approach would no longer be able to solve this. In our model, we chose to not sync service time across the materials, since this would not disturb the theoretically optimal result. In application, the model results will just be used to set new base stock levels of components. Therefore, the base stock levels of common components across the BOM can be aggregated.

5.2.2 Lead Time Customization

The Python package, Stockpyl, necessitates an integer value for process time, prompting us to designate the manufacturing process time as a single day. The presence of multiple stages may potentially result in an overall longer process duration. However, in practical scenarios, the Water Company is capable of completing the assembly within a day, assuming the absence of capacity constraints. Additionally, it should be noted that the Water Company often conducts assembly in distinct stages using a batch work approach. Consequently, the process time under such circumstances is contingent upon the specific arrangement of the batch work.

To effectively operationalize this model, it may be prudent for the Water Company to consider a more granular assumption for process time, thereby enhancing the model's real-world applicability and precision.

5.3 Future Work

In this section, we consider ways in which our model results and applicability can be improved, based on limitations discussed in section 5.2.

5.3.1 Scenario Planning for Cost Savings

In the present model, the inbound service time for materials sourced from external suppliers is considered a constant variable. In prospective research, it would be beneficial to explore the financial implications of varying this inbound service time. Once materials that significantly impact the financial performance are identified, the Water Company could strategize to reduce the inbound service time through various methods. These could include negotiating with suppliers or revising shipping methodologies.

However, it is important to note that this approach could require substantial computational resources to simulate and analyze a range of different scenarios. Therefore, the feasibility and practicality of such an approach should be assessed considering the available computational capacity.

5.3.2 Alternative Algorithms for Synchronizing Service Time Across Materials

As highlighted in Section 5.2.1, "Synchronizing Service Times Across Materials," we have approached identical materials as distinct stages in the model. A potential consolidation of these into a single stage would alter the topology of the Bill of Materials (BOM), preventing it from maintaining the structure of a spanning tree. By definition, a spanning tree comprises N stages with N-1 arcs, while the new topology could introduce a circular structure, thereby generating more than N-1 arcs.

While there is existing research addressing general networks, the application of dynamic programming becomes significantly more complex within this context. The prospect of consolidating materials and applying algorithms designed for general networks as a method to further decrease inventory costs presents a promising avenue for future exploration. It is, however, essential to consider the trade-off between potential cost savings and the increased computational complexity that such a strategy might entail.

5.3.3 Phantom Stock Removal

The repercussions of phantom stock elimination constitute a crucial yet under-investigated area within existing literature. This can be largely attributed to its distinctive correlation with the execution of multi-echelon models within the framework of a Bill of Materials (BOM). While the deployment of phantom stock is a prevalent practice across various organizations for cost accounting purposes, it poses significant obstacles for optimization endeavors. This is largely since excising these stages influences the topology of the overarching structure.

Moreover, the standard cost of phantom components frequently encompasses indirect labor costs, which might not be reflected in its child stages once the stage is removed. Nonetheless, these costs continue to be incorporated within any parent stages. This incongruity could conceivably distort the outcomes of the optimization process.

Therefore, the formulation of an effective strategy to uphold the integrity of the BOM structure, whilst simultaneously excluding phantom stages, is integral to obtaining a more precise representation of optimal inventory levels. As such, this presents a fertile field for additional scholarly investigation.

5.3.4 Impact on Supply Chain Resilience

While the primary objective function of this model is focused on minimizing holding costs, it does not account for the overall resilience of the system. Adjusting the safety stock of certain components could potentially introduce vulnerabilities in the face of supply disruptions.

These could trigger delays that ripple through the company's value chain, ultimately affecting the customer service level.

Therefore, it would be prudent to conduct an examination of the trade-off between reducing safety stock at specific stages and the prospective financial loss that could arise from a supply chain disruption at that stage. This analysis should ideally be carried out prior to implementing this model, to ensure that its application does not inadvertently undermine the system's robustness and reliability. This balancing act between cost optimization and supply chain resilience represents an important area of study for supply chain management.

6. CONCLUSION

In this capstone project, we developed a multi-echelon inventory optimization model tailored specifically for our sponsor company.

The model demonstrated the impact of inventory positioning on the overall inventory value. Notably, by consolidating approximately 67.97% of the safety stock value at the finished product level, the model presented a potential inventory allocation that would allow the company to achieve customer service levels while minimizing inventory holding costs. However, it is difficult to gauge the true financial impact of these results, because the model must be run with every SKU as an input to get a true picture of how much inventory to hold. Additionally, our sensitivity analysis revealed that varying customer service times and service levels could result in inventory cost reductions, providing managers with a flexible tool to adapt inventory policies according to changing business needs.

Nevertheless, the model has several limitations, particularly in handling identical materials as distinct stages and the challenge of synchronizing service times across materials. The model's assumption of a single manufacturing process time across components also limited its real-world applicability and precision. Consequently, the study identified potential areas of future work, including exploring the financial implications of varying inbound service time, applying other methods to consolidate materials, eliminating phantom stock, and assessing the impact of safety stock adjustments on supply chain resilience.

In summary, this work contributes to the Water Company's inventory management by providing a framework for inventory justification. It highlights the importance of multi-echelon inventory optimization models in managing complex value chains, informing strategic decisions, and enhancing financial performance. The ongoing development of these models will support adaptability and competitiveness in a dynamic business landscape.

7. REFERENCES

Badinelli, R. D. (1990). The inventory costs of common mis-specification of demand-forecasting models. The International Journal of Production Research, 28(12), 2321-2340.

Bonney, J. (2009). The great chase: Supply chain synchronicity (cover story). Journal of Commerce, 10(25), 12–16. http://0-search.ebscohost.com.ujlink.uj.ac.za/login.aspx?direct=true&db=bth&AN=42745300&site=ehost-live&scope=site

Chu, L., & Shen, Z.-J. M. (2003). Note on the complexity of the safety stock placement problem. Technical note, University of Florida.

Fisher, M. L., Raman, A., & McClelland, A. S. (2000). Rocket science retailing is almost here: Are you ready? Harvard Business Review, 78(4), 115–124. http://0-search.ebscohost.com.ujlink.uj.ac.za/login.aspx?direct=true&db=bth&AN=3261355&site=ehost-live&scope=site

Graves, S. C. (1988). Safety stocks in manufacturing systems. Journal of Manufacturing and Operations Management, 1, 67-101.

Graves, S. C., & Schoenmeyr, T. (2016). Strategic safety-stock placement in supply chains with capacity constraints. Manufacturing & Service Operations Management, 18(3), 445-460.

Graves, S. C., & Willems, S. P. (2000). Optimizing strategic safety stock placement in supply chains. Manufacturing & Service Operations Management, 2(1), 68-83.

Harris, F. (1913). How many parts to make at once. Operations Research, 38(6), 947-950.

Hoffmann, T. R., Fogarty, D. W., & Blackstone, J. H. (1991). Production & Inventory Management. South-Western Publishing Company.

Humair, S., & Willems, S. P. (2011). Technical note: Optimizing strategic safety stock placement in general acyclic networks. Operations Research, 59(3), 781-787.

Humair, S., Ruark, J. D., Tomlin, B., & Willems, S. P. (2013). Incorporating stochastic lead times into the guaranteed service model of safety stock optimization. Interfaces, 43(5), 421-434.

Inderfurth, K. (1991). Safety stock optimization in multi-stage inventory systems. International Journal of Production Economics, 24, 103-113.

Inderfurth, K., & Minner, S. (1998). Safety stocks in multi-stage inventory systems under different service measures. European Journal of Operational Research, 106, 57-73.

Jeunet, J. (2006). Demand forecast accuracy and performance of inventory policies under multi-level rolling schedule environments. International Journal of Production Economics, 103(1), 401-419.

Kimball, G. E. (1988). General principles of inventory control. Journal of Manufacturing and Operations Management, 1, 119–130.

Klosterhalfen, S., & Minner, S. (2010). Safety stock optimisation in distribution systems: A comparison of two competing approaches. International Journal of Logistics Research and Applications, 13(2), 99–120. https://doi.org/

Lopez, J. A., Mendoza, A., & Masini, J. (2013). A CLASSIC AND EFFECTIVE APPROACH TO INVENTORY MANAGEMENT. International Journal of Industrial Engineering, 20(5-6), 372-386.

Magnanti, T. L., Shen, Z.-J. M., Shu, J., Simchi-Levi, D., & Teo, C.-P. (2006). Inventory placement in acyclic supply chain networks. Operations Research Letters, 34, 228-238.

Minner, S. (1997). Dynamic programming algorithms for multi-stage safety stock optimization. Operations Research Spektrum, 19(4), 261-271.

Nassauer, S., & Terlep, S. (2022, August 28). Inventory Pileup, Uneasy Shoppers Put Retailers in Jeopardy. Wall Street Journal. https://www.wsj.com/articles/inventory-pileup-uneasy-shoppers-put-retailers-in-jeopardy-11661690106

Naughton, N. (2020). Why You Might Have Trouble Getting the Refrigerator, Can of Paint or Car You Want. Wall Street Journal. https://www.wsj.com/articles/factories-rush-to-keep-up-with-post-lockdown-shopping-11603627201

Peterson, R., Pyke, D. F., & Silver, E. A. (1998). Inventory management and production planning and scheduling. Wiley.

Rungtusanatham, M. J., Goldstein, S. M., & Schroeder, R. (2010). Operations Management: Contemporary Concepts and Cases. McGraw-Hill Education.

SAP. (n.d.). ABC/XYZ Segmentation. SAP Help Portal. Retrieved December 2, 2022, from https://help.sap.com/docs/SAP_INTEGRATED_BUSINESS_PLANNING/feae3cea3cc549aaa9d9de7d363a83e6/33e40058547f8073e10000000a441470.html

Shen, Z.-J. M. (2019). MULTIECHELON INVENTORY MODELS. In Fundamentals of Supply Chain theory. essay, John Wiley & Sons, Inc.

Silver, E. A., Pyke, D. F., & Thomas, D. J. (2016). Inventory and production management in supply chains. CRC Press.

Water Company. “Water & Wastewater Products & Services.” Water Company, 2022, https://www.Water Company.com/en-us/products--services/. Accessed 3 November 2022.

Yu, M., Tian, X., & Tao, Y. (2022). Dynamic Model Selection Based on Demand Pattern Classification in Retail Sales Forecasting. Mathematics, 10(17), 3179.