Sweet Inventory Management

An Optimization Study of Perfetti Van Melle's Inventory Policy in Vietnam

Master's Thesis
in the field of Industrial Engineering & Management

Rutger J. Habets
University of Twente
Enschede, The Netherlands

HO CHI MINH CITY

SWEET INVENTORY MANAGEMENT

AN OPTIMIZATION STUDY OF PERFETTI VAN MELLE'S INVENTORY POLICY IN VIETNAM

Author
R.J. (Rutger) Habets

A thesis submitted in partial fulfillment of the requirements for the degree of Master of Science (Industrial Engineering & Management).

Supervisors: dr. I. (Ipek) Seyran Topan - University of Twente
dr. E. (Engin) Topan - University of Twente

External supervisors: T.V.H. (Hoa) Tran - Perfetti Van Melle
M.H. (Manh Ha) Nguyen - Perfetti Van Melle

UNIVERSITY OF TWENTE, 2023

// Acknowledgments

I am thrilled to have completed my Master's in Industrial Engineering & Management, and I am grateful for the support and guidance I received along the way.

First of all, I would like to express my gratitude to Piet, Nikhil, Erwin and Tiger for giving me the unique opportunity to intern at Perfetti Van Melle, the leading company in the confectionery industry of Vietnam.

In addition, I wish to thank my company supervisors, Hoa and Manh Ha. I would like to express my appreciation to Hoa for her unwavering support and attention to detail, which were pivotal to the success of my thesis project. I am grateful to Manh Ha for his consistent help in advancing the project's progress and for his valuable guidance, both personally and professionally. Furthermore, I would like to thank all my colleagues for the fun team trips marked by a lot of laughter and fantastic memories that we made. I cherish the beauty of your culture and country, which will always hold a special place in my heart.

I am recognizing the efforts of my university supervisors, Ipek and Engin. Thank you, Ipek, for taking me as a graduate student and for your continuous support over the years. I have benefited greatly from your meticulous editing and from our extensive evaluation sessions, which I enjoyed every moment of. I also wish to express my gratefulness to Engin for his encouragement and valuable insights into the theoretical aspects of my project.

Then, I would like to extend my gratitude to my parents, sister and friends as they have provided immense support to me during my academic journey.

In closing, I am grateful to everyone who supported me throughout this experience, and I am excited to see what the future holds. With the completion of this thesis, I have reached a significant milestone and I am eager to put my knowledge and skills into practice in the business world.

For now, I would like to wish you an interesting read.

Best regards,

Rutger J. Habets
Ho Chi Minh City, Vietnam

Executive Summary

In today's fast-paced and competitive global economy, effective supply chain management is pivotal for increasing organizational effectiveness and profitability. One critical aspect of supply chain management is *inventory management*, which involves the overall process of ensuring the availability of products. It plays a major role in balancing costs and service goals. This thesis considers the inventory policy of the *Perfetti Van Melle* (PVM) in Vietnam, a privately owned international company active in the confectionery industry, offering popular brands such as *Chupa Chups*, *Mentos* and *Alpenliebe*.

With a north-to-south distance of about 1,650 km, Vietnam is commonly divided into three regions: the North, Center and South. In pursuit of a high service level, PVM operates three Distribution Centers (DCs), one in each region. The DCs are monitored independently but offer lateral transshipment, as the southern DC can move excess inventory to the other two DCs. Overproduction, resulting from the *minimum production quantity* (MPQ), leads to a *partial pooling* effect, where inventory exceeding a threshold is distributed among the other DCs.

The stock levels of the DCs are managed by a *periodic review policy* with a variable lot size, also known as an (R,S) replenishment policy, where a production order is placed when the inventory position is lower or equal to the order-up-to-level S in every review period R. PVM applies this replenishment policy to its entire product portfolio, consisting of three different product types, of which this thesis focuses on the XX core products. To date, there is no tailored approach per product (class), i.e., R is XX days and S is XX days of stock on hand for all products.

The Vietnamese confectionery industry is subject to two types of customers, which cover almost the entire sales volume of PVM. *Traditional trade* (TT) represents XX.X% of the market, the *modern trade* (MT) channel makes up XX.X% and other channels cover the remaining XX.X%. PVM gives order priority to MT over TT, but this is primarily done on an *ad hoc* basis, without a specified strategy to handle it.

From the above introduction, it is clear that PVM's inventory management relies on practical experience rather than a theoretical foundation. This has led to the implementation of one generic inventory strategy for all products and customers, resulting in undesirable inventory levels. These inventory levels create an imbalance in stock, with excessive stock causing unnecessary holding costs but primarily shortages leading to stockout costs and dissatisfied customers.

The objective of this thesis is to explore the opportunities for optimization of PVM's inventory policy based on the trade-off between inventory costs and fill rates. As such, the central research question is: How do we improve Perfetti Van Melle Vietnam's inventory policy to maximize product availability and minimize inventory costs?

Product classification is a commonly used strategy to simplify inventory management by setting stock control parameters and fill rates per class rather than for each product separately. This thesis applied the ABC and ABC-XYZ classifications. The latter is a combination of the ABC classification (based on annual turnover) and the XYZ classification (according to demand fluctuation).

The order-up-to-level is the main control parameter for PVM's (R,S) policy. It is the maximum inventory position and equals the expected demand during the review period and lead time, plus the safety stock. The safety stock plays an essential role in preventing shortages, but balancing it against holding costs is challenging.

Next, we examined the partial pooling strategy and considered the possibility of implementing a complete pooling approach. The complete pooling approach involved using the entire stock available at the SDC for transshipments.

To deal with customer differentiation, we extended the traditional (R, S) model by introducing a critical level k as a second control parameter. This led to an (R, S, k) inventory policy, where low-priority TT customers are not served once the inventory falls below the critical level k. According to the literature, the critical level policy should reduce the overall fill rate but increase the fill rates for the MT priority class.

Determining the optimal inventory parameters for a real-life (R, S) policy is difficult, particularly when incorporating the additional critical level parameter k and pooling effect. This thesis used the speed of analytical computations and the accuracy of simulation-based optimization to arrive at an effective solution. The analytical approximations served as inputs for the simulation, which were then further optimized with a greedy algorithm. The algorithm is designed to make a locally optimal choice in each stage, with the goal of achieving a suitable solution.

We created a conceptual model to mimic the inventory system, which formed the basis for developing the simulation model in *Python*. We conducted 3 experiments, which include the ABC-XYZ classification, ABC classification and complete pooling. We found that the partial pooling policy without critical levels was the best choice. Despite the critical levels slightly enhancing MT performance, they led to increased overall costs or unmet target fill rates. Furthermore, both the ABC and ABC-XYZ classifications showed similar performance improvements. Therefore, we suggest integrating the ABC classification given its ease of implementation. We have found the following order-up-to-levels:

Table 1 Improved order-up-to-levels (in days of stock on hand).

Class	A	B	C
Southern DC	XX	XX	XX
Northern DC	XX	XX	XX
Central DC	XX	XX	XX

The control parameter settings resulted in the performance of Tables 2 and 3, excluding external influences like supply disruptions. From these tables, it is visible that the overall performance increased by 7.7% and inventory costs were reduced by VND[1] XX,XXX million for the system without external constraints. Assuming that a state without disruptions is realized XX% of the time, annual savings would be VND XX,XXX million, equivalent to approximately € XX,XXX.

Table 2 Comparison of fill rates between the current and improved policy.

Policy	Total	Traditional Trade	Modern Trade
Current	XX.X%	XX.X%	XX.X%
Improved	XX.X%	XX.X%	XX.X%

[1] At the time of writing, the exchange rate stands at about VND 25,400 for € 1.

Table 3 Comparison of costs (in millions) between the current and improved policy.

Policy	Inventory costs	Holding costs	Stockout costs
Current	VND XX,XXX	VND XX,XXX	VND XX,XXX
Improved	VND XX,XXX	VND XX,XXX	VND XX,XXX
Difference	- VND XX,XXX	+ VND XX,XXX	- VND XX,XXX

The main advice for PVM is to implement the proposed replenishment settings in its inventory management. Additionally, the supply planning department should revise the ABC(-XYZ) classification monthly and run the simulation every quarter year to incorporate new data and products. In future, the simulation could be implemented in PVM's SAP data system. During this thesis research, another area for improvement identified was PVM's data quality. To achieve practical and cost-efficient data quality, the company should focus on building advanced data structures and investing in automation tools that aid in the data cleansing process.

In this thesis, we made valuable contributions to PVM's inventory management, including a stockout decision-making tool, a mother-mapping lookup document that connects data across departments, and a template for the ABC-XYZ classification. We also developed a simulation tool with an implementation manual to ensure a sustainable solution for the company. Additionally, our thesis adds to the existing literature by considering both customer differentiation, transshipments and substitution in the framework of multi-products, multi-locations problems with product classification for fast-moving products. We combined simple analytics, simulation, and a straightforward greedy algorithm to fill this research gap. Our findings indicate that in our dynamic and fast-moving case study context with strict fill rate requirements, the implementation of the critical level policy did not result in any improvement. Furthermore, we have observed that the existing partial pooling strategy performs effectively by efficiently managing the surplus inventory generated by the minimum production quantity.

This thesis has limitations concerning the supply constraints, demand distributions and data sample used. We optimized the unconstrainted performance, but in reality, supply chain issues prevent the unconstrainted state from being realized for a considerable fraction of the time. On the other hand, the system's flexibility, through *ad hoc* decision-making, is greater than what could have been modeled. The thesis is only applicable to products with normally distributed demands, resembling fast-paced behavior, and the data sample used was limited to three years. With an average solution error of 7.27%, the proposed algorithm leads to sufficient solutions, but better alternatives could be explored. Future research should also further explore the system's resilience and consider the potential of safety stock to respond to sudden supply chain disruptions, particularly with the increasing frequency of natural disasters, catastrophes, and pandemics.

Keywords: inventory management, (R,S) policy, fast-moving products, product classification, multi-locations, customer differentiation, partial pooling, simulation, greedy algorithm.

Table of Contents

List of Acronyms

DC	Distribution Center, where SDC, CDC, NDC are the Southern, Central and Northern DCs, respectively.
IBP	Integrated Business Planning
MPQ	Minimum Production Quantity
MT	Modern Trade
PVM	Perfetti Van Melle (Vietnam)
S&OP	Sales and Operations Planning
SCM	Supply Chain Management
SQ	Sub-Question
TT	Traditional Trade
VND	Vietnamese Dong

List of Symbols

Indices

d	distribution center $d \in D$
i	product $i \in I$
j	product class $j \in J$
n	simulation run $n \in N$
t	time $t \in T$

Inventory Management

IL	inventory level
IP	inventory position
k	critical level
k_{ss}	safety factor
L	lead time
mu (μ)	demand
R	review period
S	order-up-to-level
ss	safety stock

Performance Measures

$beta$ (β)	fill rate
eta (η)	fill rate ratio
h	unit holding cost
p	unit price
s	shortage cost
TC	total inventory costs
v	annual value

Statistics

CI	confidence interval
CV	coefficient of variance
$sigma$ (σ)	standard deviation
$\bar{X}$	mean

1 Introduction

This section serves as a global introduction to the research performed at the case company Perfetti Van Melle (PVM). The thesis examines the current inventory management of its subsidiary in Vietnam by focusing on the replenishment policy. Section 1.1 motivates the relevance of this research from a general perspective, and Section 1.2 offers background information on Perfetti van Melle (Vietnam). Section 1.3 explains the problem statement, and Section 1.4 defines the research setup and methodological framework. Section 1.5 presents the intended research deliverables.

1.1 Significance

In times of globalization and soaring competitiveness, *supply chain management* (SCM) has become a key strategic element for increasing organizational effectiveness and profitability (Crittenden, Crittenden, & Crittenden, 2017; Garcia, Marchetta, Camargo, Morel, & Forradellas, 2012). As defined by the Council of *Supply Chain Management Professionals* (CSCMP) (Vitasek, 2013), a leading-edge organization in the field of SCM, it covers the planning and management of all activities involved in sourcing, procurement, conversion, and all logistics management activities to fulfill customer requirements. So, a proper SCM strategy enhances beneficial results in terms of process efficiency, inventory levels, customer satisfaction, quality, costs, and delivery (Núñez-Merino, Maqueira-Marín, Moyano-Fuentes, & Martínez-Jurado, 2020).

Inventory management, described by the CSCMP as the overall process of ensuring the availability of products (Vitasek, 2013), has a significant impact on the orchestration of a company's supply chain activities (Gümüs & Güneri, 2007). It plays a major role in balancing costs and service goals, and acts as a critical buffer against demand and supply uncertainty (Gümüs & Güneri, 2007). Inventory usually represents 20 to 60% of a manufacturing company's total assets. Since holding inventory entails high logistics costs, the benefits from inventory must outweigh the expenses. Inventory optimization covers the trade-off between these elements by minimizing the stock in the entire supply chain while ensuring product availability.

Inventory management has been re-recognized as a strategic component of the national economy in emerging markets such as Vietnam (Hai Nam, 2019). To date it is accounting for a large proportion of enterprise costs, making it a critical barrier to the competitiveness of enterprises. This thesis reviews the inventory policy of the company Perfetti Van Melle Vietnam.

1.2 Introduction to Perfetti Van Melle

Perfetti Van Melle (PVM) is a privately owned and integrated international group active in the confectionery industry. The company was established in 2001 through a merger of *Perfetti S.p.A.* (founded in 1946) and *Van Melle NV* (1900). It is one of the world's largest manufacturers and distributors of candies and chewing gums in over 150 countries worldwide (Perfetti Van Melle, 2019). With a team of more than 17,900 employees, PVM posted a net turnover of € 2,672 million in 2021. The sales volumes included 405,908 tons kg, which for example weighs the equivalent of around 67,651 adult elephants.

PVM offers strong brands, such as *Mentos*, *Chupa Chups* and *Alpenliebe* (see Fig. 1.1). Apart from these global brands, the company matches the tastes of individuals in local and regional markets by offering numerous tailored-made products like *Fruittella*, *Frisk*, *Smint*, *Stimorol*, *Look-O-Look*, *Golia* and *Big Babol*. The ability to create innovative brands, varieties and flavors that meet the different requirements and opportunities of local markets is one of its keys to success.

Fig. 1.1 Global brands of PVM.

1.2.1 Perfetti Van Melle Vietnam: Introduction

Worth knowing, Vietnam has only been open to business with the rest of the world for a couple of decades (Wunker, 1994). Following this move to an economic policy of accepting and encouraging foreign investment in 1986, its market became an exciting, albeit difficult, place to do business.

In 1997, PVM established a subsidiary in Vietnam to expand its business into a market with a strongly increasing young population. The subsidiary was set up under a joint venture agreement with Saigon Foods Stuff Company, in which Perfetti S.p.A. contributed 70% of the capital. A partner was selected due to the government's force as well as its partner's capabilities in producing glucose - a critical raw material for Perfetti.

In the following years, Vietnam carried out great efforts to improve the legal environment to turn Vietnam into an even more attractive country for investing (Vo & Nguyen, 2012). The new governing law strengthened the role of the private sector, e.g., by allowing private enterprises to trade goods (e.g., glucose) without import and export licenses. Consequently, PVM switched its ownership to 100% in 2002 to increase growth opportunities. The enterprise still maintains a good relationship with its former partner up to now. Appendix A.1 presents a detailed infographic overview of PVM Vietnam's milestones over the past 25 years.

Today *PVM Vietnam* is much ahead of its competitors and leads the confectionery segment in Vietnam through its Chupa Chups, Alpenliebe and Golia - occupying about XX.X% of the candy market share in 2021. Its production plant, as seen in Fig. 1.2, processes an impressive XX,XXX tons kg of candy annually – half of which is designated for domestic consumption. PVM Vietnam has a team of over 2,000 employees, of which its 700 sales representatives deliver the demand to 150,000 stores.

Fig. 1.2 Bình Dương factory.

1.2.2 Perfetti Van Melle Vietnam: Plan and Deliver

From this point on PVM refers to Perfetti Van Melle Vietnam. PVM's *Plan and Deliver* team ensures that the right products are produced at the right quantity and time to meet sales' requirements while balancing production capacity, with optimum material levels, and preparing sufficient raw and packaging materials for production.

Mrs. Viet Hoa, the supervisor of this research project, leads the Plan & Deliver team, of which Fig. 1.3 shows the organizational chart. This project is performed for the Supply Planning department (marked in green) and focusses on domestic trade. Supply Planning focuses on inventory management to fulfill the demand plan while meeting the financial and service goals of the company.

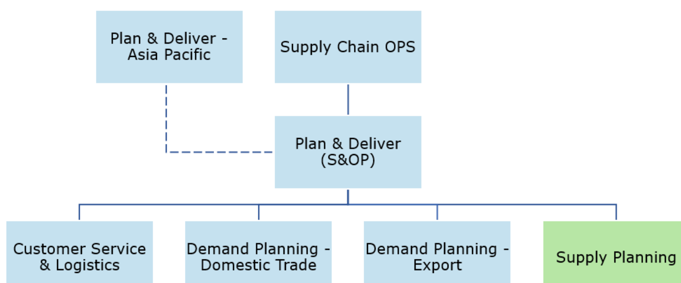

Fig. 1.3 Organizational chart of Perfetti Van Melle Vietnam: Supply Chain OPS - Plan & Deliver.

1.2.3 Perfetti Van Melle Vietnam: Domestic Trade

Two types of sales channels cover almost the entire sales volume of PVM, namely traditional trade (TT) and modern trade (MT). TT is a complex distribution network of small retailers, kiosks, hawkers, wholesalers and distributors. It builds on interpersonal relations between customers and retailers and is the main source of sales for PVM, accounting for XX.X% of their total value. Traditional retail is prone to erratic demand, which may lead to empty shelves or the need to push alternate products onto the customers. Fig. 1.4 below gives an idea of what a typical TT retail store looks like.

On the other hand, MT operations are more planned and consistent with inventory management. These outlets are chains or groups of businesses, such as supermarkets, mini-markets and hypermarkets. This channel generates XX.X% of PVM's total revenue and is expanding rapidly.

Fig. 1.4 Traditional trade store in Ho Chi Minh City.

Then, the Vietnamese confectionery market is influenced by seasonality. In general, the industry has two seasons: high season and low season. The period from February until August is defined as the low season. The high season starts in September when demand surges due to the Halloween and Christmas holidays. Sales usually peak in December and January, just before the *Tet Holiday* (traditional Vietnamese Lunar New Year). During the high season, competition is fierce with seasonal players targeting the market.

Fig. 1.5 presents an overview of the supply chain processes, which can be explained using the metaphor of a river. Starting at the upstream end of the chain, manufacturing (in the factory) and temporary storage (in a small adjacent warehouse) are the first steps. Next, primary transportation takes place, moving the stock to the Distribution Centers (DCs): South, Central and North. Inventory management is crucial at this stage, ensuring that each of these DCs has enough stock to meet their respective demands.

The South Distribution Center (SDC) has a two-fold purpose: (i) to keep enough inventory to serve the southern demand and (ii) to act as a consolidation center for the northern and central regions (see dashed arrows in Fig. 1.5). The consolidation stock occurs as a result of overproduction, which in its turn is a result of the Minimum Production Quantity (MPQ) set by the factory. The SDC will separate the stock based on its intended use. The North and Central Distribution Centers (NDC, CDC) only need to keep enough inventory to meet their demands. This thesis targets the inventory management of stock and therefore focuses on the corresponding tactical-level processes in the outbound logistics network.

Downstream the inventory management processes, secondary transportation takes place. This is the transfer of stock to agents and wholesalers. These agents and wholesalers are third-party providers that buy stock from PVM. Finally, sales representatives (employed by PVM) transfer the stock to their customers by truck or motorcycle.

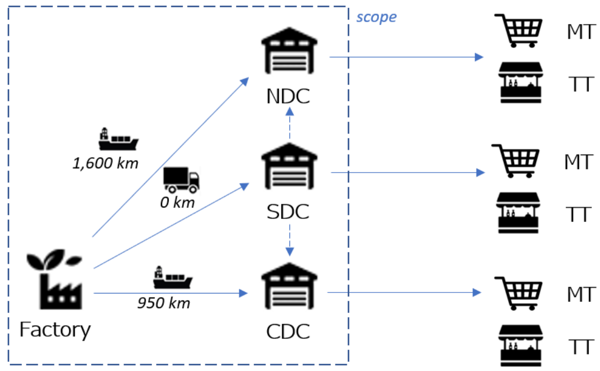

Fig. 1.5 *Schematic overview of the logistics processes in domestic trade.*

1.3 Problem Statement

As introduced in Section 1.1, it is important to realize we are living in a dynamic and competitive global economy. Fast and massive changes are affecting businesses and impacting traditional channels of distribution (Princes, 2020; Vo-Thai, Hong-Hue, & Tran, 2021). Industries are evolving with many reaching maturity and searching for ways to create growth. Companies are focusing on strategic renewal, searching for creative solutions to improve quality, reduce costs, raise the bar on customer service, manage risk, and increase efficiency. They need a strategy for managing all resources that go toward meeting customer satisfaction. So, supply chain and logistics capabilities, as well as competitive strategies have become an important field in business improvement

PVM also realizes that supply chain logistics has become a unique element and a source of international competitiveness. For many years, the company maintained a competitive rate in which logistics costs (i.e., the costs for transportation and storage) were only about XX.X% of the revenue. This figure is lower than the average 5-10% logistics costs reported by the Ministry of Industry and Trade in Vietnam, highlighting PVM's significant advantage over other candy manufacturers and distributors. Though, to remain a competitive advantage, PVM must strive to continuously improve supply chain performance and inventory management.

In addition to low costs, the company strives for a more resilient and agile supply chain to maintain its position as the market leader. An agile supply chain strategy requires the chain to react quickly and flexibly to changing demand, making a high service level the market winner (Jammernegg & Reiner, 2007). The *service level* is a metric, shown in percentages, which captures the ability to satisfy demand or responsiveness (Vitasek, 2013). Having high service levels is key in meeting all possible inventory demand requirements, ergo, it plays a major role in keeping customers satisfied in the uncertain logistical environment.

Safety stock is seen as an important inventory driver due to its critical role in mitigating the risk of getting out of stock (Rădăşanu, 2016). Its level must be high enough to cover fluctuations in demand, inaccurate forecasts, shortages in supply and insufficient production capacities. But having too much safety stock may result in high holding costs. This brings forth the dilemma between efficiency (minimizing inventory costs) and responsiveness (maximizing product availability).

Advanced data analytics and automation can help to effectively manage safety stock, allowing organizations to better understand their performance and identify areas of improvement. Despite being aware of this, PVM faces difficulties in consolidating and analyzing its stock data to gain valuable insights. Currently, daily reviews are conducted, but the data is neither stored nor properly analyzed, resulting in limited visibility into the performance.

Additionally, its inventory management is primarily based on experience rather than a theoretical SCM foundation. This has led to the same inventory strategies being applied across all of their products and customer channels, resulting in an untailored approach and inventory levels that are undesirable. Poor inventory levels create an imbalance in stock, where excessive stock results in unnecessary holding costs and obsoletes. On the other hand, shortages lead to stockouts, which reduce the overall service level and profit due to lost sales and customers, and relatively expensive emergency transfers.

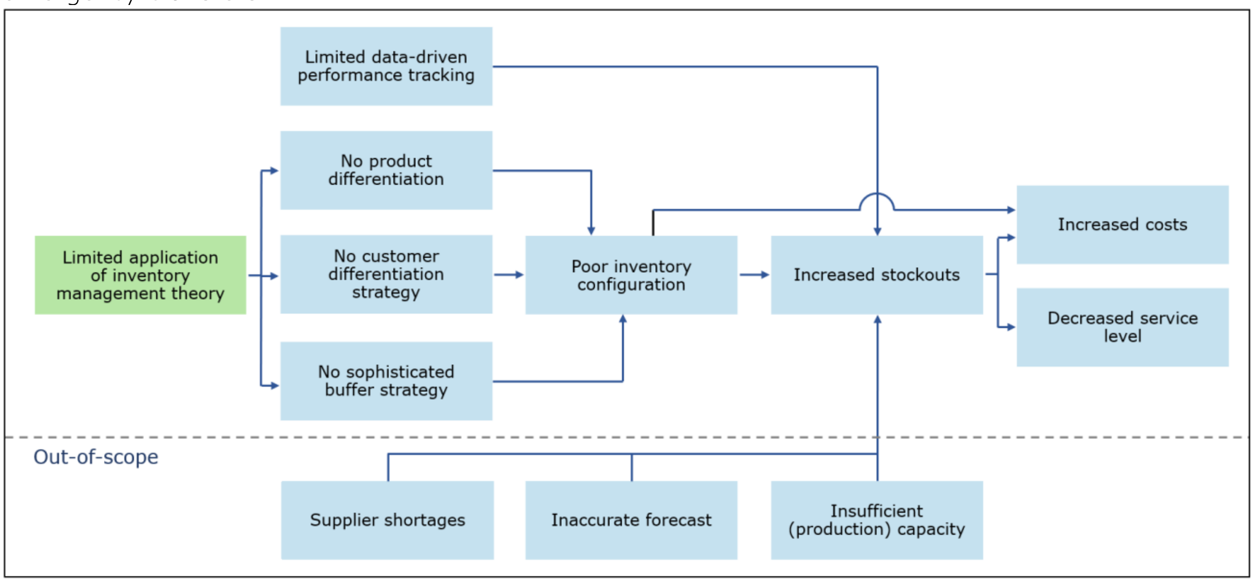

Fig. 1.6 Problem cluster including core problem (green).

It is clear from the problem cluster in Fig. 1.6 that there is significant room for improvement in PVM's inventory policy. The objective of this thesis is to get a deeper understanding of the opportunities for optimization through a thorough (data) investigation, considering trade-offs among the inventory costs and service level. With this in mind, the thesis aims to answer the following central research question:

How do we improve Perfetti Van Melle Vietnam's inventory policy to maximize product availability and minimize inventory costs?

1.4 Research Design & Methodology

The central research question provided the starting point and path for all valuable knowledge development by setting the shape, direction and progress of the thesis. It can be described as the structural axis of the investigation. As shown in Table 1.1, the research question was divided into four sub-questions (SQs) to guide the collection of missing knowledge, while also establishing the outline of the thesis.

Table 1.1 Subdividing the central research question.

Central Research Question: *"How do we optimize Perfetti Van Melle Vietnam's inventory policy to maximize product availability and minimize inventory costs?"*	
SQ1 Context Analysis	*What does the current inventory policy look like and what is the performance? (Section 3)*
SQ2 Literature Study	*What methods for inventory policy optimization are suggested in the literature? (Section 2)*
SQ3 Solution Design	*What is the best way to model the optimization model? (Section 4)*
SQ4 Conclusion	*What is the recommended inventory policy configuration? (Sections 5 and 6)*

SQ1 *What does the current inventory policy look like?*

It was essential to get a grasp of the current inventory policy, meaning that the related processes, activities, stakeholders, customers, and key performance measures were reviewed. We acquired information through (informal) interviews, documentation and observations at PVM. Operators, logistics analysts and managers were interviewed because of their expected knowledge of the company. A data analysis then provided us with deeper insights into the inventory control's performance, limitations and core problem. We extracted the data from PVM's information system SAP and prepared it in Excel and Power BI. Together with stakeholders, we analyzed the data on the source, patterns and modifications and performed extensive cleansing to guarantee credibility.

SQ2 *What opportunities and methods for optimization of the inventory policy are suggested in the literature?*

We performed a literature review to define approaches for finding the optimal configuration of the inventory policy. To find similarities and gaps between the research problem and similar theoretical cases, scientific articles and books were searched in the online databases *Scopus*, *Web of Science* and *Google Scholar*. The search query considered the context keywords *safety stock*, *reorder* point, *order-up-to-level*, *periodic review*, *inventory policy* and *inventory control* in combination with the stockout keywords *lost sales* and *backorder*, as well as with the methodological keywords *multi-criteria*, *decision model*, *decision tool*, *probabilistic*, *stochastic*, *heuristic*, *model*, *optimization*, *simulation* and *algorithm*. We used various combinations and synonyms of the keywords to get a large selection of relevant literature. We also excluded the keywords *multi-echelon*, *two-echelon, spare parts* and *queueing* from the search strings, as these words led to articles that were irrelevant to this thesis. We further broadened our search by looking into *product classification*, *lateral transshipment*, *substitutes* and *customer differentiation*, as these topics provided us with extra areas for improvement. We restricted our search to peer-reviewed scientific journals written in English, as most of the high-quality research is typically published in journals.

When identifying an article with relevance to this thesis, we also reviewed the publication's reference list of relevant publications for further research. The literature study first identified the key drivers in inventory policy related to product availability and costs, with a strong focus on the safety stock, order-up-to-level and periodic review. Then, this literature study elaborated on the concepts and methods that are known for examining and enhancing inventory policies. The literature was also consulted to acquire information about an optimization tool, a simulation that employs a *greedy algorithm*, which served as the foundation of the optimization.

SQ3 *What is the best way to model the optimization method?*

We designed a discrete event simulation model to determine near-optimal inventory policy settings. The model can be considered as convincing proof of the findings from the system analysis and literary theory. We determined the initial decision variables by analytical approximation, then optimized them using a greedy algorithm. We conducted experiments and performed a sensitive analysis to ensure the robustness of the model. We validated the optimization model by consulting both logistical stakeholders and programming experts, extensively tracing and debugging, using historical data, scenario testing, and comparing the model behavior with the real world and literature.

SQ4 What is the recommended safety stock configuration to increase the performance of the inventory policy?

Based on the findings of sub-question 3, we presented the most appropriate configuration to improve the inventory policy. Ultimately, the thesis concluded the research and provided recommendations regarding the improvements and opportunities for further research.

To summarize the research design, Fig. 1.7 depicts a brief overview of the methodology that was used to answer the research questions. This flowchart serves as a "skeleton", around which information was gathered, required to answer the central research question.

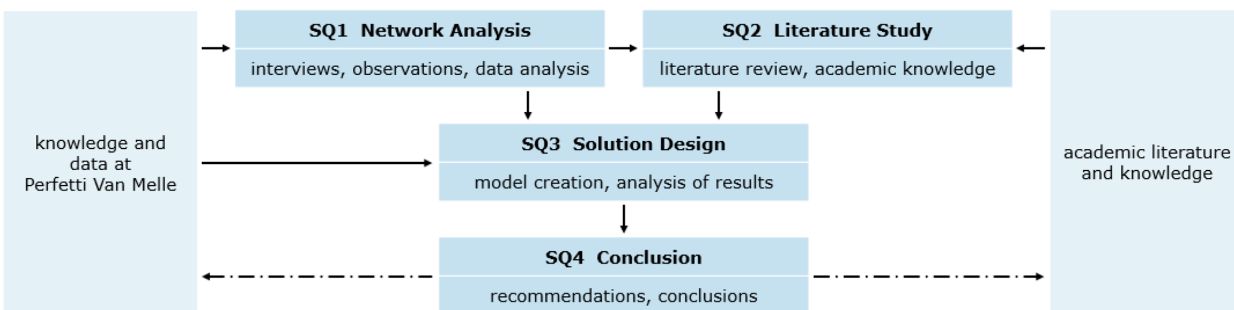

Fig. 1.7 Flow diagram of the research structure.

1.5 Deliverables

Upon completion of this research, the following deliverables had been presented:

- advice on the appropriate inventory policy with recommended settings for the safety stock, order-up-to-level, product classification and customer classification;
- a reusable optimization tool that enhances decision-making in the inventory policy;
- notable suggestions for further research.

2 Literature Study

In our literature review, we aimed to identify the optimal configuration of an inventory policy and to compare it with similar theoretical cases. In Section 2.1, we examine the drivers of inventory, and Section 2.2 outlines product classification. Section 2.3 evaluates the four standard replenishment policies, while Section 2.4 outlines how to configure the replenishment policy corresponding to PVM. As a means of improving efficiency, Sections 2.5 and 2.6 introduce transshipments and customer differentiation, respectively. Section 2.7 recommends appropriate optimization methods, and Section 2.8 wraps up the literature review in the context of PVM.

2.1 Inventory Drivers

This thesis aims to maximize product availability while minimizing inventory costs. This trade-off requires a proper understanding of the factors that drive an inventory. A large amount of research has been done about (qualitative) inventory drivers, including the paper by Chinello, Lee Herbert-Hansen and Khalid (2020). The authors reviewed 56 articles and proposed a decision-making framework of 10 essential optimization drivers in inventory management. Table 2.1 depicts the framework.

In accordance with the supply planning stakeholders, we classified the drivers into three types based on their relevance to this research. A low relevance means that the driver was out-of-scope for this research. A medium relevance implies that the driver is considered but not directly affected by this research, i.e., the driver is out of control. A driver with high relevance was directly subject to and in control of this research.

The inventory drivers are easy to understand from the description in Table 2.1. For this thesis, general insights are sufficient for the low and medium-relevant drivers. These drivers are not investigated. However, the prominent research drivers are discussed thoroughly. These are product classification and the replenishment policy.

In the “Impact of Inventory” column, the arrows indicate the effect of increasing a specific driver, unless otherwise stated (e.g., the product life cycle length is decreased). The forecast accuracy is an example of an increased driver that positively impacts service levels, as depicted by the upward-pointing arrow.

Table 2.1 *Inventory drivers framework (Chinello et al., 2020).*

Driver	Description	Impact on Inventory	Relevance
Forecast accuracy	Degree of the closeness of demand forecasted to actual demand	↑ Service level ↓ Stock levels	Medium
Demand variability	Changes in demand from period to period	↑ Risk of stockout, safety stock, forecast complexity, setup and ordering cost	Medium
Supply variability	Changes in supply volumes or lead times	↑ Safety stock, stock levels	Low
Product classification	Classification of products based on business impact	↑ Control focus, service level ↑↓ Safety stock	High
Product life cycle length	Duration of the sequence of product introduction to market, growth, maturity, and decline	Short life cycle: ↓ Predictability of stock needed ↑ Forecasting complexity	Low
Product variety	Number of different products in the product portfolio	↑ Difficulty in fulfilling demand, complexity, stock levels, inventory, and backorder costs ↓ Fill rates	Medium
Production process and capacity	Steps taken and techniques adopted to transform inputs into outputs	Production process simplifies and capacity increases: ↓ Raw materials, work in progress ↑ Finished goods, capability to meet target customer service level	Low
Replenishment policy	Reordering or replenishment process defines the review period for reordering and the order quantity	Small order size: ↓ Setup time, holding costs, stock levels ↑ SC flexibility, ordering costs	High
Shipping frequency	Frequency of shipments to replenish inventories	↑ Transportation costs ↓ Order size, inventory costs, stock levels	Medium
Lead time	The time between the placement of an order and delivery.	Short lead times: ↓ Inventory levels, stockout risk, order size	Medium

2.2 Product Classification Methods

Product classification is widely adopted by organizations because it simplifies the task of inventory management by setting stock control parameters and service levels per class rather than for each product individually (Teunter, Babai, & Syntetos, 2010). There are numerous types of classifications, but we are limited to the most straightforward and renowned. This decision was taken in consultation with the management, as product classification is a new concept for PVM. This leads us to three very well-known classifications, namely ABC, XYZ and combined ABC-XYZ.

2.2.1 ABC Classification

ABC inventory classification systems are widely used to streamline the organization and management of inventories consisting of many products (Stojanović & Regodić, 2017). The *ABC classification* distinguishes products by ranking them according to the annual value v_i, which is given by:

$$v_i = p_i * \mu_i \tag{2.1}$$

with unit price p_i and annual demand μ_i. Based on revenue, the ABC method divides the amount of control effort per product. Products are grouped according to the Pareto principle, also known as the 80:20 power law, which is based on the observation that there is a small number of products that contributes dominantly to the achieved sales results. This results in the forming of three product classes:

A-class This group covers about 20% of the portfolio and is responsible for 75-80% of the total revenue. A-class products are the most important because of their large contribution to the total revenue. Therefore, these products need to be controlled tightly and monitored closely;

B-class Class B products cover the next 10-15% of the products in the network and account for 15-20% of the total company result. B-products are less important, due to their smaller contribution to the total sales, and (at least) require computed-based inventory management;

C-class These are the last 50% of products in the system, accounting for roughly 5% of the total revenue. Generally, class C products are slow movers and therefore face a higher obsolescence risk. Inventory management of these products should be kept as simple as possible (Stojanović & Regodić, 2017).

Fig. 2.1 shows a graphical representation of the ABC method. It is important to mention that the ABC distribution is arbitrary, and the groups are defined in accordance with the company's needs. So the analysis should not be based solely on the product's contribution to the total revenue because products with a small contribution, but essential to the business, can also be seen as A-products (Stojanović & Regodić, 2017).

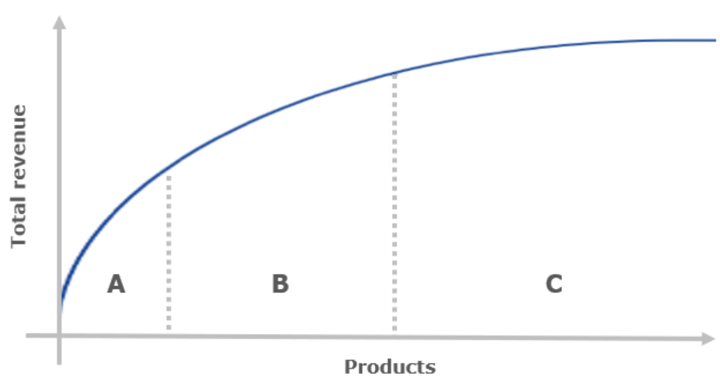

Fig. 2.1 Graphical representation of the ABC classification.

The ABC classification is known for its simplicity, but the same has been criticized for using only one criterion for classification. In 1987, Flores and Whybark were the first who recommended using a multi-criteria classification with the traditional ABC method as the primary criterion. Nowadays, it is generally accepted that the ABC analysis should include several criteria (Silver, Pyke, & Thomas, 2017; Stojanović & Regodić, 2017). This brings us to the XYZ classification.

2.2.2 XYZ Classification

In the *XYZ classification*, the ranking is done by the criterion of the demand variability compared to the average demand. In other words, the products are classified based on the fluctuation in demand, also known as demand uncertainty. The demand uncertainty is determined with the help of the coefficient of variation (CV). The CV_i is calculated with eq. 2.2, where σ_i is the demand standard deviation and $\bar{X}_\mu$ the demand average (Stojanović & Regodić, 2017).

$$CV_i = \frac{\sigma_i}{\bar{X}_\mu} \tag{2.2}$$

Products can be assigned to one of the following three categories:

X-class	This class (0-10% of the product portfolio) experiences a continuous demand, characterized by very slight fluctuations, making this group the most promising for forecasting. X-products normally have a CV of less than 0.5;
Y-class	These products (10-25%) have some fluctuations in demand and are harder to forecast. The CV is usually between 0.5 and 1.0;
Z-class	This group (usually 25% or more) faces the most fluctuations in demand, for which there are usually hardly predictable causal factors. The Z-class also encompasses low-selling products, characterized by high variable demand, making it difficult to get a reliable demand forecast. These products have a CV greater than 1.0 (Scholz-Reiter, Heger, Meinecke, & Bergmann, 2012; Stojanović & Regodić, 2017).

Fig. 2.2 visualizes typical movement patterns of demand of the X-, Y-, and Z-groups. Similar to the ABC distribution, ranks are arbitrary for the XYZ analysis.

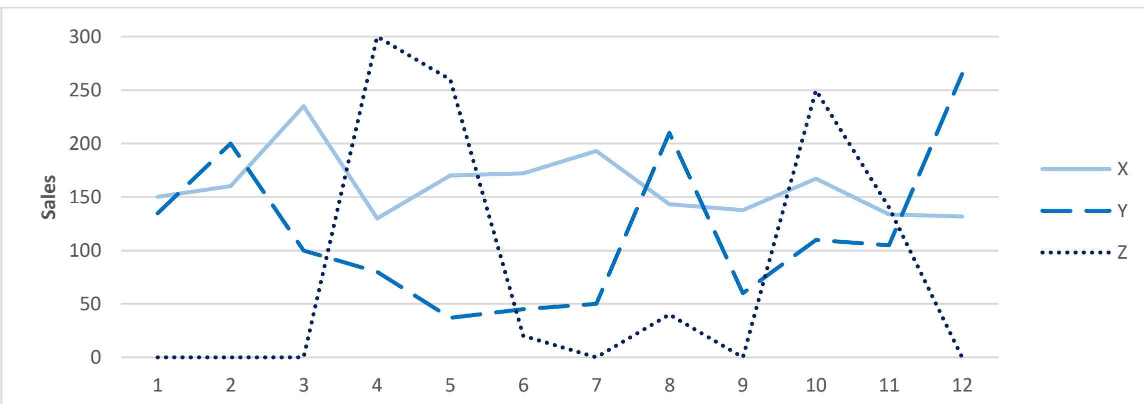

Fig. 2.2 Typical demand patterns for the XYZ classification.

2.2.3 ABC-XYZ Classification

As explained, it is interesting to not only group products by their importance or complexity but to integrate both methods. The renowned *ABC-XYZ classification* divides the products into nine categories based on their revenue contribution and demand uncertainty.

Table 2.2 gives us a matrix with the main characteristics per product class. Products present in class AX are products that are most important and are easy to forecast, whereas products present in product group CZ are the least important, but are the most difficult to forecast. The benefit of determining inventory policies per class becomes clear now, as the number of integrated policies shrinks from possible hundreds (or more) to a maximum of nine.

Table 2.2 ABC-XYZ classification.

	X	Y	Z
A	High-value percentage Continuous demand High predictive value	High-Value percentage Fluctuating demand Medium predictive value	High-Value percentage Irregular demand Low predictive value
B	Average-value percentage Continuous demand High predictive value	Average-value percentage Fluctuating demand Medium predictive value	Average-value percentage Irregular demand Low predictive value
C	Low-value percentage Continuous demand High predictive value	Low-value percentage Fluctuating demand Medium predictive value	Low-value percentage Irregular demand Low predictive value

2.3 Replenishment Policy

In inventory management, it is essential to determine the appropriate timing and quantity of a reorder. These decisions impact the cycle and safety inventories as well as the fill rate (Chopra & Meindl, 2016). To make informed decisions, it is important to consider the *inventory position* (IP), calculated as the inventory level minus outstanding production orders (Axsäter, 2015).

Inventory policies are generally classified into four categories, as shown in Table 2.3. Inventory policies are categorized as continuous or periodic review. In *continuous review*, the inventory is continuously tracked and an order is placed when the IP drops below the reorder point. In *periodic review*, the IP is only monitored at certain intervals, and an order is placed when the IP is equal to or lower than the reorder point at the beginning of the review period. Additionally, the policies can be categorized into fixed or variable lot sizes. *Fixed lot size* policies have a constant order quantity (or a multiplicity of the order quantity), while *variable lot size* policies have a variable quantity and aim to reach a certain IP (Silver *et al.*, 2017).

Table 2.3 Inventory control policies.

	Continuous review	Period review
Fixed lot size	(s, Q) or (s, nQ)	(R, s, Q) or (R, s, nQ)
Variable lot size	(s, S)	(R, S) or (R, s, S)

(s, Q) or (s, nQ) policy

The (s, Q) policy is a policy with continuous review and fixed lot size. When the IP drops below the *reorder point* s, a production order is placed. The reorder point is equal to the expected lead time L during demand plus the safety stock. The policy becomes (s, nQ) when it is also possible to order a multitude of the *lot size* Q (Silver *et al.*, 2017).

(s, S) policy

The (s, S) policy is a policy with continuous review and variable lot size. Similar to the (s, Q) policy, when the IP drops below the reorder point s, a production order is placed. The difference is that order quantity is fixed in the (s, S) policy. It equals the difference between the *order-up-to-level* S and the IP (Bartmann & Bach, 2012).

(R, s, Q) or (R, s, nQ) policy

The (R, s, Q) policy is a policy with a periodic review and fixed lot size. A production order will be created when IP is lower or equal to s at the start of the review period R. It is also possible to purchase an integral multiple of the order quantity Q such that IP is raised to a value between s and $s + Q$, then the policy is defined as (R, s, nQ) (Janssen, Heuts, & De Kok, 1998).

(R,S) or (R,s,S) policy

The (R,S) and (R,s,S) policies are policies with a periodic review and variable lot size. In the (R,S) policy, IP is raised to S in every R. In the (R,s,S) policy, an order is placed when IP drops to or below s. The order quantity equals the difference between S and IP (Silver *et al.,* 2017). Fig. 2.3 illustrates the (R,S) replenishment policy.

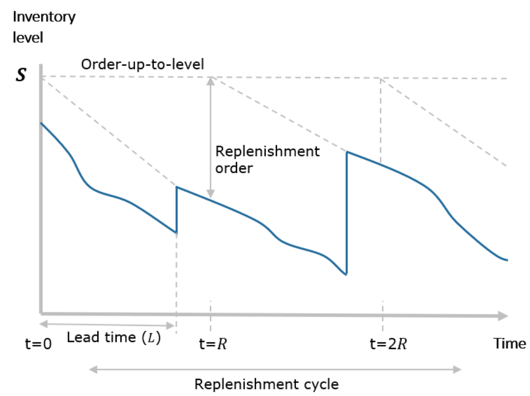

Fig. 2.3 Order-up-to-level with periodic review.

2.4 Replenishment Parameters

PVM follows an (R,S) policy for its replenishment cycle. This section elaborates on S and other criteria involved. R is pre-determined by the company and therefore not addressed in this literature review. Before discussing the parameters of the replenishment policy, it is important to familiarize oneself with the safety stock and undershoot.

2.4.1 Safety Stock

It is known that no supply chain can operate without safety stocks, which are required to manage supply and demand variability (Gonçalves, Sameiro Carvalho, & Cortez, 2020). The standard deviation of the forecast error under R and L is used to calculate the *safety stock* using the equation:

$$ss = \sigma_{R+L} k_{ss}, \tag{2.3}$$

where k_{ss} is the safety factor for demand variation (Silver *et al.*, 2017).

Maintaining a high level of service is crucial for customer loyalty. Many companies aim for a 95% service level, which is a balance between opportunity and operational costs. However, determining the optimal service level for each product can be challenging and requires a thorough understanding of the product and industry. An ideal safety stock strategy should be small enough to reduce inventory-related costs while satisfying demand on time and ensuring a high service level. Customer behavior to stockouts varies by product, so it is important to consider this when setting service levels. To simplify the process, the ABC-XYZ analysis can be used to assign specific service levels to each category (Rădăşanu, 2016).

2.4.2 Order-up-to-Level

The order-up-to-level S is the maximum IP and equals the expected demand during $R + L$, plus ss (Van der Heijden, 2020). The equation for S under periodic review is:

$$S = \hat{\mu}_{R+L} + ss, \tag{2.4}$$

where $\hat{\mu}$ represents the expected demand during $R + L$ (Van der Heijden, 2020). If IP falls below S during review, a replenishment order will be initiated. The order quantity is determined by subtracting IP from the S.

2.5 Transshipments Methods

A renowned strategy for minimizing costs while maintaining an adequate level of inventory is *lateral transshipment*, which involves "sharing or rotating stock in a multi-location inventory system at the same echelon in the supply chain" (Nakandala, Lau, & Shum, 2017).

There is a large body of research on transshipments (Agrawal & Smith, 2015), of which a significant stream addresses the inventory control of spare parts, such as the work of Kranenburg and van Houtum (2009). The authors develop a method using the Erlang loss model to analyze multi-item, multi-location systems. But their assumption that demand follows the Poisson distribution, which is common in literature on spare parts, is not applicable to systems with high demand (Escalona, Ordóñez, & Kauak, 2017). Systems with high demand typically have normally distributed demand, so our review excludes the literature in the context of spare parts.

Instead, we are interested in a periodic review system with transshipments, multi-items with a high demand per period, and multi-locations, which is only addressed by a few inventory models (Hu, Watson, & Schneider, 2005).

Nakandala *et al.* (2017) classify transshipments as either proactive or reactive. *Proactive transshipments* are preventative and involve balancing inventory levels to reduce the risk of future shortages. They are beneficial when inventory handling costs are high (Kumari, Wijayanayake, & Niwunhella, 2022). *Reactive transshipments*, also called emergency transshipments, are triggered by an actual inventory shortage after demand has been realized. In such cases, it is cost-effective to implement a reactive strategy if the expenses of the transshipments and the inability to immediately fulfill customer demand are less than the costs of maintaining high inventory levels.

Paterson, Teunter and Glazebrook (2012) introduce a hybrid approach that combines the benefits of both transshipment policies. Rather than merely replenishing an inventory shortage, their policy sees each reactive transshipment as an opportunity for proactive stock redistribution by transferring extra inventory, incorporating a proactive element.

Another classification criterion for transshipment is the pooling method, which is either complete or partial. In *complete pooling*, the transshipping location shares all its available stocks, whereas *partial pooling* only agrees to transship while its inventory level is above a fixed threshold. This threshold can be determined as a percentage of the current inventory, a fixed quantity of inventory, or a function of the remaining time until the next replenishment (Kumari *et al.*, 2022). Fig. 2.4 shows partial pooling in the replenishment policy resulting from the MPQ, with pooling occurring at inventory positions above S. This scenario is relevant to PVM.

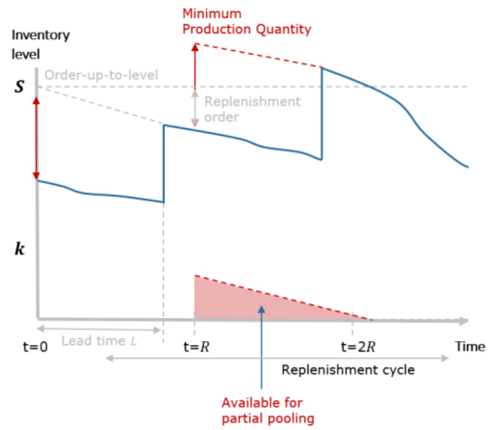

Fig. 2.4 *The (R,S) policy with partial pooling as a consequence of the MPQ.*

Additionally, the transshipment problem can be categorized as centralized or decentralized. A *centralized problem* aims to minimize the overall cost across all locations. In a *decentralized problem*, each location endeavors to maximize its own profitability by deciding on the quantity and price of items to be transshipped (Kumari *et al.*, 2022).

2.6 Customer Differentiation Methods

PVM has segmented its customers into two demand classes: (i) the high-priority MT class that requires high service levels and (ii) the low-priority TT class that represents the retailers that have to settle for a lower service level.

Despite a vast body of literature on *customer differentiation* in inventory management, comparatively fewer studies have focused on inventory systems with lost sales. This was pointed out by Paul and Rajendran (2011) and still holds today.

One of the pioneering studies in this area of lost sales was conducted by Cohen, Kleindorfer and Lee (1988). They were the first to develop a lost-sales model for two customer classes, using an (R, s, S) inventory system. Under this policy, high-priority demand is satisfied first every R, while low-priority demand is met with the remaining inventory.

Tempelmeier (2006) clearly explains the renowned approaches for inventory management with respect to differing service requirements of its customers: (i) Use a single stockpile with a safety stock that meets the highest service level requirement among all customers, also referred to as the round-up policy. This leads to excess inventory for customers with lower service requirements than the maximum; (ii) A separate stock policy, where customers are served from a common stockpile and customer-specific safety stocks are set similar to the round-up policy, this also leads to excess stock and does not utilize the possible inventory pooling effects; (iii) implement a *rationing policy* that, when inventory has fallen below a critical level k, stops serving low-priority customers from the common stockpile.

In this thesis, we focus on the last option, which is known as the *critical level policy*. It can be implemented for several ordering and review policies. For example, the traditional (R, S) model is extended using a critical level policy to an $(R, S, \boldsymbol{k})$ inventory model, where k denotes the critical level for the low-priority customer class, expressed as an percentage of S. Fig. 2.5 visualizes the critical level policy as part of the replenishment cycle.

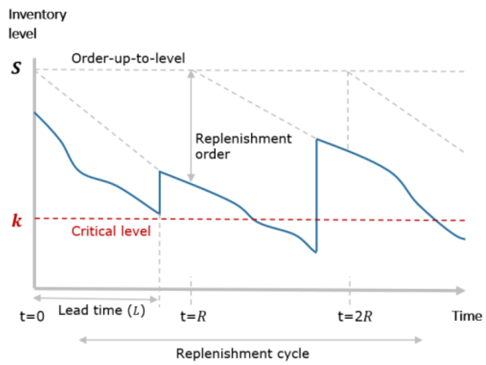

Fig. 2.5 (R,S,k) critical level policy.

The critical level policy, first coined by Veinott Jr. (1965) is an efficient way of providing differentiated service levels by rationing inventory among different customer classes. It enables the pooling effect while also providing different levels of service without the need to maintain a large inventory for classes that require less service level than the maximum (Escalona *et al.*, 2017). The critical level policy results in a lower overall fill rate but higher fill rates for high-priority classes (Nahmias & Demmy, 1981).

Despite being a longstanding approach for dealing with customer differentiation, critical level policies in high-demand systems have not received much attention. Escalona *et al.* (2017) were the first researchers to examine a critical level policy using a continuous distribution to model demand for fast-moving items. They formulated a linear problem to determine the policy parameters, with a continuous review period and back ordering. Escalona, Angulo, Weston, Stegmaier and Kauak (2019) later extended this research by studying the effects of two service-level measures on the design of the critical level policy, and using a global search algorithm to find the optimal solution.

2.7 Optimization Methods

Finding the right inventory parameters for an (R, S) policy in PVM's complex distribution network is extremely challenging due to the presence of correlations in demand and supply among regions. There are several ways to find these parameters, including analytical approximation and simulation.

Analytical approximation can provide approximate solutions quickly, and with relatively little computation, can be useful for understanding the behavior of the system . However, the dynamic nature of inventory is the major obstacle for inventory control practitioners and makes most analytical models either over simplistic or computational intractable. To overcome this limitation of existing analytical models, the literature suggest optimization via simulation because of its capability of handling variability (Chiadamrong & Piyathanavong, 2017; Silver *et al.*, 2017).

Optimization via simulation can be useful in cases where the supply chain is highly complex or dynamic, and the relationships between the various factors are not well understood. In these cases, simulation can provide a more detailed and nuanced view of the supply chain, allowing for more accurate optimization of the inventory parameters. But it can be computationally intensive and may not provide accurate results if the simulation does not accurately reflect the real-world supply chain (Chiadamrong & Piyathanavong, 2017; Göçken & Dosdoğru, 2015).

When deciding whether to use optimization via simulation or analytical methods, it is important to consider the complexity and dynamic nature of the supply chain. For simpler supply chains, analytical methods may be sufficient, while for more complex supply chains, simulation may be necessary. In some cases, it may also be beneficial to use a combination of both approaches, incorporating the insights from analytical approximation into the simulation to improve the accuracy of the results (Chiadamrong & Piyathanavong, 2017).

Analytical models can be used to analyze and optimize the performance of simple transshipment models, but these models often rely on restrictive assumptions like those involving instantaneous replenishment and negligible transshipment lead time (Tagaras, 1999). Relaxing these assumptions can complicate the mathematical analysis, as it requires considering the interrelationships between demands, transshipment quantities, and in-transit stock, implying that the state space should be expanded (Tlili, Moalla, & Campagne, 2012). It can also have secondary effects on future replenishments and the establishment of safety stocks. Simulation, on the other hand, is a powerful tool for generating numerical representations of complex transshipment systems. It allows researchers to analyze the interactions and causal relationships among system components and evaluate the performance of various competing models (Nakandala *et al.*, 2017; Tlili *et al.*, 2012). Despite the consensus among logistics experts that simulation, potentially paired with heuristics, should be promoted as a supplement to traditional analytical methods, there is currently a limited number of published simulation studies in the transshipment literature (Nakandala et al., 2017).

Dealing with customer differentiation (i.e., critical level) problems in periodic review inventory systems is complicated, making it very difficult to solve these problems using analytical methods or mathematical programing. Instead, the literature suggests that suitable approaches are the use of simulation or heuristic algorithms, sometimes combined. Studies considering simulation by (Azimi, Ghanbari, & Mohammadi, 2012; Li & Chen, 2010) and heuristic algorithms by (Escalona *et al.*, 2019; Najjartabar Bisheh, Nasiri, Esmaeili, Davoudpour, & Chang, 2022; Paul & Rajendran, 2011) are examples of this.

A group of well-known heuristics that can be applied are greedy algorithms. These algorithms always make the decision that seems best at the current moment. That is, it makes a locally optimal choice in each stage with the hope of obtaining a globally optimal solution as a result. Despite its effectiveness for a vast array of problems, it should be noted that greedy algorithms do not guarantee the optimal solution (Cormen, Leiserson, Rivest, & Stein, 2007).

A classic greedy algorithm consists of five elements: a candidate set, selection function, feasibility function, objective function, and solution function. The candidate set serves as the pool of potential solutions from which the algorithm will choose, while the selection function selects the best candidate to be added to the solution. The feasibility function checks if a candidate can be used to contribute to the solution, and the objective function assigns a value to the solution or partial solution for evaluation and comparison. Finally, the solution function indicates when a complete solution has been found (Cormen *et al.*, 2007).

The greedy algorithm is often a suitable choice for optimization simulations, as it can rapidly provide a satisfactory solution (Rai, Gross, & Ettam, 2015). The globally optimal solution may not always be discovered, but the greedy heuristic is often sufficient in delivering an adequate solution in a reasonable amount of time.

2.8 Conclusion of the Literature Study

Classification of products is widely adopted by organizations to simplify the task of inventory management, by setting stock control parameters and service levels per class rather than for each product separately. It has been decided, together with the management, to limit the possible classifications to convenient and renowned ones. This leads us to the ABC and ABC-XYZ classifications. The latter is a combination of the ABC classification (grouping products according to their annual turnover) and the XYZ classification (grouping based on demand fluctuation).

PVM adopted an (R, S) replenishment policy for its inventory management, which has a periodic review and variable lot size. In this thesis, review period R is out-of-scope and set at XX days by the company. A production order will be created when the inventory position IP is lower or equal to the order-up-to-level S at the start of the review period.

Transshipments, or the sharing of inventory across multiple locations in the supply chain, are a cost-efficient method of managing inventory levels. Our review evaluates three criteria for transshipment classification, leading us to classify the PVM transshipment policy at the southern DC as a partial pooling strategy, with a centralized costs review and decentralized fill rate evaluation. PVM's transshipment policy adopts a reactive approach, similar to the hybrid policy developed by Paterson and colleagues in 2012. During a stockout, it permits the transfer of additional inventory, thereby incorporating a proactive element.

PVM has segmented its customer channels into two classes, namely the high-priority MT and the low-priority TT class. In this thesis, we focus on a critical level policy to deal with customer differentiation, also referred to as a rationing policy. A critical level policy ensures that, when inventory has fallen below a critical level k, it stops serving low-priority customers from the common stockpile. The traditional (R, S) model is extended using a critical level policy to an (R, S, k) inventory model. By implementing this modified policy, it becomes possible to offer varying levels of service without the need to maintain a large inventory for customer classes that require less service than the maximum level. The critical level policy results in a lower overall fill rate but higher fill rates for high-priority classes.

Finding the correct inventory parameters for a real-life (R, S) policy is challenging. There are several methods for determining these parameters, including analytical approximation and simulation. Analytical approximation can offer approximate solutions quickly and with minimal computational effort, and it can be useful for understanding the behavior of the system. However, the difficult nature of PVM's inventory system makes it hard to study using analytical methods. To overcome this limitation, the literature suggests optimization via simulation because of its capability of handling variability. Therefore, we will utilize simulation to ensure a more accurate optimization of the inventory parameters. It is also a powerful tool for generating numerical representations of complex transshipment systems and is promoted as a supplement to traditional analytical methods by many logistics experts. Regarding customer differentiation, the literature suggests that suitable approaches are the use of simulation, heuristic algorithms, or a combination of both.

Well-known heuristics that can be applied with simulation are the greedy algorithms, as they can rapidly provide an adequate solution for many problems. These algorithms make a locally optimal choice in each stage with the hope of obtaining a globally optimal solution. To the best of our knowledge, the application of greedy heuristics in a high-demand periodic review system with customer differentiation and transshipments has not been studied thoroughly, offering the potential for new discoveries through this thesis.

3 Context Analysis

The purpose of this section is to gain a thorough understanding of the problem by examining the current inventory system. It has been divided into three sub-sections: Section 3.1 delineates the processes and stakeholders involved with the inventory policy. Section 3.2 analyzes the performance of key drivers associated with the inventory policy. Section 3.3 concludes the major findings of this context analysis.

3.1 Field of Research

This sub-section elaborates on the stock replenishment activities of PVM. Section 3.1.1 introduces the relevant inventory drivers for this thesis, while Section 3.1.2 discusses the outline of the company's logistics network. Section 3.1.3 provides insights into the integrated business planning cycle, and Section 3.1.4 explains the replenishment cycle itself. Section 3.1.5 portrays the differences in the behavior of customer channels, and Section 3.1.6 covers the different product types of PVM.

3.1.1 Inventory Drivers

Section 2.1 introduced the most common drivers used in inventory management. In accordance with the supply planning stakeholders, we classified the relevance of these drivers for this thesis into three categories: low, medium and high. Drivers with a low relevance were beyond the scope of this thesis. Medium relevance means that the driver is relevant to the thesis, but not directly influenced by it, i.e., the driver is out of control. A driver with high relevance was directly subject to and in control of this thesis.

Table 3.1 lists the relevance for the inventory drivers in scope. From this table, it is visible that the product classification and replenishment policy are the main inventory drivers, while the forecast accuracy, demand variability, product variety and lead time were also in the research scope but out of control. This section covers all inventory drivers of medium to high relevance, excluding product classification which can be found in the solution in Section 4.1. The arrows in the "Impact of Inventory" column represent the impact of increasing a particular driver, unless noted differently.

Table 3.1 Inventory drivers framework.

Driver	Description	Impact on Inventory	Relevance
Forecast accuracy	Degree of the closeness of demand forecasted to actual demand	↑ Service level ↓ Stock levels	Medium *Section 3.2.3*
Demand variability	Changes in demand from period to period	↑ Risk of stockout, safety stock, forecast complexity, setup and ordering cost	Medium *Section 3.2.2*
Product classification	Classification of products based on business impact	↑ Control focus, service level ↑↓ Safety stock	High *Section 4.1*
Product variety	Number of different products in the product portfolio	↑ Difficulty in fulfilling demand, complexity, stock levels, inventory, and backorder costs ↓ Fill rates	Medium *Section 3.1*
Replenishment policy	Reordering or replenishment process defines the review period for reordering and the order quantity	Small order size: ↓ Setup time, holding costs, stock levels ↑ SC flexibility, ordering costs	High *Section 3.1.4*
Lead time	The time between the placement of an order and delivery.	Short lead times: ↓ Inventory levels, stockout risk, order size	Medium *Section 3.1.2*

3.1.2 Logistics Vietnam

The research considers the logistics and inventory management of PVM in Vietnam. Vietnam is a long, narrow country, resembling the letter 'S' as it snakes down the South China Sea from north to south. The capital city, Hanoi, is located in the northern region, while the country's largest city, Ho Chi Minh City, is in the south. The country has a north-to-south distance of around 1,650 km and is typically divided into three main areas: the North, Center and South.

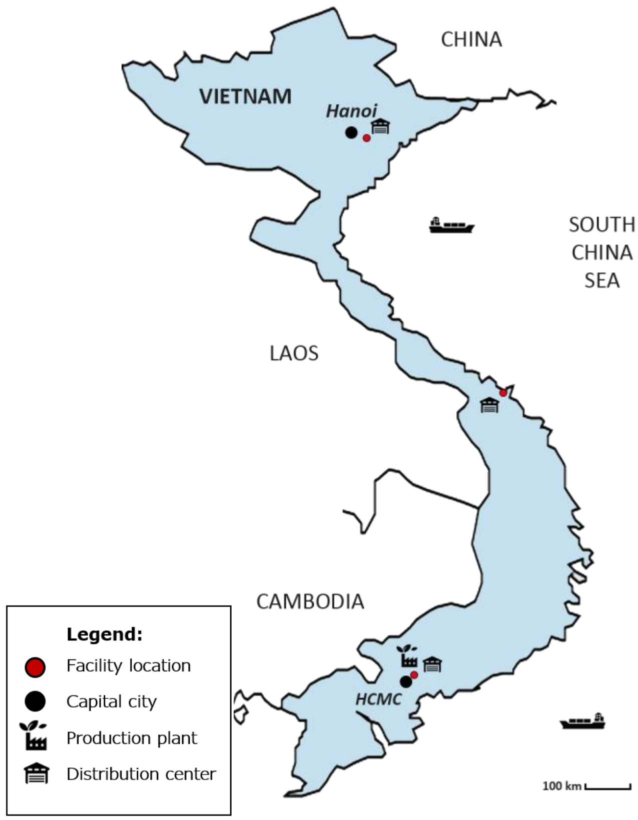

Fig. 3.1 PVM's facility locations in Vietnam.

Distribution Centers

In pursuit of a high service level, PVM has three DCs, owned by third-party providers, covering each region and offering exclusive storage space to the company on a fee-per-pallet basis (see Fig. 3.1). Multiple DCs enable the company to react quickly and with greater flexibility (Christopher, 2000). It also enhances the risk-diversification effect by spreading the impact of stock-outs, including risks associated with disruptions (Schmitt, Sun, Snyder, & Shen, 2015).

The company's factory is located in Bình Dương province in southern Vietnam. The SDC is adjacent to the factory and has two purposes:

(i). to keep inventory for the demand of southern Vietnam, while also serving as a safety stock buffer station;

(ii). to serve as a consolidation center for the northern and central regions, as the factory (warehouse) is only capable of storing limited quantities of stock. Excess inventory from PVM's production, due to Minimum Production Quantity (MPQ), is stored at the SDC as a buffer for lateral transshipments following a partial pooling strategy. The transshipment policy is reactive, with a proactive element for stockouts, so the policy allows the transfer of additional inventory in case of stockouts. Fig. 3.2 illustrates the corresponding transportation flows of PVM's logistics network, including transshipment flows shown as dashed lines. The factory is modeled as an ample supplier.

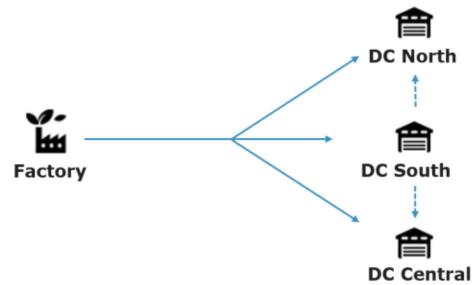

Fig. 3.2 *Supply flow chart of PVM's logistics network.*

Transportation Modes

Transportation between the factory and DCs is performed by third-party logistics providers and coordinated by the logistics department. Regular transportation is usually by sea, with rail transportation used for emergencies to avoid stockouts or eliminate them quickly. Although emergency transportation is faster than regular shipments, it is also more expensive and does not include the option for free storage. Furthermore, emergency transportations are typically arranged on an *ad hoc* basis in response to sudden stockouts and relies on real-time insights to utilize potential opportunities, such as adding low-stock products or upcoming promotions to the emergency shipments. Lateral transshipments from the SDC also enhance the distribution network's agility and resilience to stockouts. Table 3.2 presents the lead times per transportation mode.

Table 3.2 *Transportation lead times.*

From	To	Lead time sea	Lead time rail
Factory	SDC	-	-
Factory / SDC	NDC	XX days	XX days
Factory / SDC	CDC	XX days	XX days

3.1.3 Integrated Business Planning

In order to gain a thorough understanding of the replenishment and distribution procedures, it is key to explain PVM's *Integrated Business Planning* (IBP). IBP is a management approach for translating desired business outcomes into commercial, financial and operational resource requirements, with the overarching goal of maximizing profit while minimizing risk. It is used for *sales and operations planning* (S&OP), over a 12-month rolling horizon, and has a monthly cycle of re-planning that enables executives to make decisions based on the recommendations provided.

PVM has implemented S&OP as it ensures: (i). inventory reduction, (ii). inventory control, (iii). accountabilities, (iv). one integrated plan, (v). reduction in obsolescence, (vi). improved inventory turns, and (vii). inventory and cost visibilities. Fig. 3.3 below visualizes the corresponding IBP cycle.

Fig. 3.3 *The IBP cycle.*

The IBP cycle starts with the statistical forecast on day XX. The demand planner provides a 12-month statistical forecast by the *Arkieva* planning software. In general, two years of historical data serve as input for the IBP forecast, but five years of data are available if necessary.

The portfolio review (day XX) is the first official meeting of the cycle. This meeting captures the risk and opportunities and incorporates the predicted growth rate and available budget of PVM. Marketing acts as the owner of this meeting and aims to create a balanced portfolio that meets the requirements of the category and product families for each of the markets and channels. Promotional products, deletions, and new product developments for the planning horizon are revised or included. The initial forecast might be adjusted as a result of decisions made during this meeting.

The demand review takes place on day XX. The objective is to create a valid and realistic annual forecast for the market that reflects the consensus of sales and marketing. The forecast includes projected volumes, EBIT (earnings before interest and taxes) and product family contribution. The portfolio pipeline and deletions are important inputs. Margins and forecast volumes are updated for approved portfolio changes, brand market plans, channels and distribution plans, price changes, and distributor and customer services. In the end, it is essential that budget gaps are recognized and understood, risks managed and opportunities are listed, and an action plan to clear finished good surpluses is established. This meeting is the last phase of the demand side in S&OP, which means that the final forecast will be delivered to the supply planners afterward.

The supply planner tries to realize the demand plan while meeting the company's financial and service goals. Supply planning factors in all aspects related to inventory, production and logistics. It endeavors to meet the demand forecasts as efficiently as possible. Supply planning also focuses on lead times, minimum order quantities, leveling production, and managing safety stocks. This is because a supply plan captures the demand plan and all component data, generating a master production schedule. Once in place, the supply planning examines capacity and its impact on resources, while making revisions as required in the process.

On day XX, the supply plan is provided in the supply review. The supply plan, which is created based on the unconstrained forecast, is balanced with the available capacity and the finished goods inventory policy. This meeting examines the operational performance, unit plans and assumptions made, while related recovery recommendations are drafted. The outcomes are an approved supply chain capacity plan, alignment on exceptions to the finished goods inventory policy, clear insights into the exceptions of costs of goods sold, and listed risks and opportunities.

The objective of the finance review (day XX) is to present the updated volume and revenue forecast for the business as input into cash flow forecasts. The meeting delivers an approved product family and market forecast, recommendations and scenarios for the next IBP cycle, managed costs of goods sold exceptions and managed risks and opportunities. This meeting can be considered the final rehearsal for the executive meeting.

The cycle is concluded with the executive review on day XX. This review instantly aims to yield a balanced volume plan for the business unit while taking into account the capacity and inventory policies and moving ultimately to a projected profit and loss statement for the business unit across the IBP planning horizon.

3.1.4 Replenishment Policy

In the second week of the S&OP cycle, supply planning receives the demand forecast from demand planning. This forecast serves as input for monitoring the current inventory levels (e.g., by tracking the days of stock on hand) and weekly *master production schedule* that outlines what products are needed, in what quantity, and when. To determine the feasibility of this production plan, the production levels are matched with the available capacity through a *rough-cut capacity plan*. Unfortunately, there is no historical data available on the capacity being constraining for production. So, for convenience of the thesis, it has been assumed that the production capacity is ample for the market demand. But, in reality, the capacity is estimated to be sufficient for more than XX% of the time.

Supply planning reviews the inventory level every XX days, and places a production order if the inventory position is equal to or lower than the order-up-to-level. If the position is above the order-up-to-level, nothing is done until at least the next review moment. So the company uses a periodic review policy, with a review period R of XX days and an order-up-to-level S of XX inventory days on hand for its entire portfolio. The order quantity equals the difference between the inventory position and order-up-to-level, which means that the order quantity is variable. Every production order must be large enough to guarantee cost-efficient production, which is ensured by a *Minimum Production Quantity* (MPQ). The MPQ is calculated by the factory and is out of the thesis scope. When we compare the replenishment policy described above to the findings from the literature study (see Section 2.3), it is clear that supply planning applies an (R, S) policy. This policy is also configured with the company's enterprise resource planning system *SAP*.

3.1.5 Customer Behavior in Case of Stockouts

PVM's domestic trade can be divided into modern trade (MT) and traditional trade (TT). Both channels are served by a make-to-stock production strategy but are managed separately given their different buying behavior and market fragmentation (see 1.2.3):

Traditional Trade (XX%) stores place their orders via selected sales agents or mega wholesalers. When the demand exceeds the inventory, the customer facing an out-of-stock situation chooses a *substitute product* if an alternative product from the same product type is available for purchase.

Whenever no similar products are available, a random substitute product or (partially) *lost sale* is the outcome. Lost sales are missed selling opportunities for products that are out-of-stock.

Modern Trade (XX%) involves a more planned and organized approach to distribution and logistics management. MT customers directly send their orders to PVM's DCs, where the *customer service and logistics* department processes the order. MT has order priority over TT when customers from both channels order the same product. This is a result of:

(i). PVM's business strategy, where privilege is given to the smaller sales channel;
(ii). MT customer behavior, as orders are always (partially) dropped whenever a product is out-of-stock. They do not accept substitutes and may even leave the company because of dissatisfaction.

3.1.6 Product Types

The goal of this research is to identify the best inventory control policy per product per DC. In the ideal situation, all inventory control policies are investigated per product, to see which one performs best for each product. However, PVM Vietnam has numerous different products, so this would be very time-consuming. For practical reasons, we applied the classification to group products.

Before deep-diving into the product classification, which is part of the solution of this thesis (see Section 4.1), it is important to understand PVM's different product types, which are:

(i). *core products*. These are the company's year-round products that are part of its portfolio for more than one year. To date, PVM has a total of XX core products;
(ii). seasonal products. These are linked to a specific time of the year (e.g., Christmas or Tet holiday) and are unavailable the rest of the year. PVM releases seasonal products to create some buzz around its brands, which is accompanied by promotions to surprise and excite customers;
(iii). new product developments. Innovation plays an essential role in the competitive confectionery industry. New product developments are launched to cultivate, maintain and increase the company's market share by satisfying customer demand. New product developments are evaluated after one year and will be accepted in PVM's core product portfolio if the performance is desired or deleted otherwise.

This thesis focuses on the replenishment policy of XX core products at PVM and does not include seasonal products or new product developments.

3.2 Data Analysis

This section analyzes the data to get a deeper understanding of the performance of PVM's inventory management. First, Section 3.2.1 explains the key performance measures used in this thesis. Section 3.2.2 examines the demand trends and Section 3.2.3 analyzes the forecast performance of PVM. Based on a VBA tool, the current inventory management performance is estimated in Section 3.2.4. Finally, Section 3.2.5 describes the methods used for cleaning and validating the data.

3.2.1 Performance Measures

Performance measures are quantifiable metrics, that are of strategic importance, to determine a company's progress toward a specific goal. In this thesis, the main performance measures are the (order) fill rate and total (logistics) cost.

The percentage of (order) fill rate is the fraction of customer demand that is met routinely, i.e., without lost sales (Silver *et al.*, 2017), which is a common performance measure for DCs. It is closely tied to customer satisfaction, as satisfied customers see their orders being fulfilled immediately by available stock. The fill rate is given by eq. 3.1. To exemplify, assume that 100 candies are ordered of which 95 are satisfied. This leads us to a fill rate of 95%.

$$fill\ rate = \frac{total\ orders\ shipped\ (in\ demand\ volume)}{total\ orders\ placed\ (in\ demand\ volume)} * 100\% \tag{3.1}$$

The second performance measure used is the total inventory cost. The total inventory cost TC is a product of the holding cost and stockout cost. The holding cost consists of pallet space charges for holding on-hand inventory. The cost of running out-of-stock is associated with the inability to satisfy demand, i.e., the lost sales. Eq. 3.2 provides the corresponding cost function.

$$TC = holding\ cost + stockout\ cost \tag{3.2}$$

This cost construction is not used by PVM currently, as the company has not defined the cost associated with getting out-of-stock because it is very difficult to measure the lost sales. Together with the commercial finance and supply planning departments assumptions on lost sales were made in this thesis (see Section 3.1.5 for the assumptions).

3.2.2 Demand Analysis

The core inventory drivers of this thesis are demand variability and forecast accuracy. The demand was analyzed per sales channel on variability and seasonality. This analysis was also conducted for every sales region (south, central, north) because there are three DCs. The forecast accuracy per channel was analyzed and quantified.

This thesis uses the available demand data from 2019 to 2021. In its most general form, the systematic component of demand data contains a level, a trend, and a seasonal factor. Overall, PVM's market share is increasing every year and its sales volume also experiences a growth trend. However, due to the big impact of COVID-19, its market share still increased in 2021 while the sales volumes were reduced. This decreasing demand is visualized in Fig. 3.4, which shows the demand for PVM from February 2019 until January 2021. For 2022, the growth rate (or trend) is expected to be between XX% and XX% compared to 2021.

Fig. 3.4 Domestic demand data (2019-2021).

In order to determine the seasonal factors, the demand was divided into 12 time periods of one month. For every period, the average demand was divided by the average of the total annual demand. Therefore, every seasonal factor lies between zero and one and represents the fraction of the total demand assigned to each month.

From Fig. 3.5 it can be clearly seen that all sales regions have similar behavior in demand for TT. Peak season starts in September, which is in line with PVM's business strategy, and ends in January. As indicated by the legend, the SDC and CDC are combined. This was necessary because accurately splitting the channel demand for these channels proved to be complex because both regions were previously served from the south.

Fig. 3.5 Monthly demand data (2019-2021) of traditional trade.

The same analysis leads us to Fig. 3.6 for MT, which again shows a similar behavior among the regions. Although the seasonal effect is more visible for the southern and central regions, compared to the northern region, it was decided to take the average to determine the seasonal factors. This decision was made for simplicity.

Fig. 3.6 Monthly demand data (2019-2021) of modern trade.

Table 3.3 below provides the average seasonal indices per sales channel. Traditional trade and MT follow the same seasonal patterns but TT faces a steeper decline in February while MT's demand volume gradually decreases after peak season. The trend during the high season is comparable.

Table 3.3 Seasonal factors for traditional trade and modern trade (2019-2021).

Month	Traditional Trade	Modern Trade	Month	Traditional Trade	Modern Trade
January	X.XX	X.XX	July	X.XX	X.XX
February	X.XX	X.XX	August	X.XX	X.XX
March	X.XX	X.XX	September	X.XX	X.XX
April	X.XX	X.XX	October	X.XX	X.XX
May	X.XX	X.XX	November	X.XX	X.XX
June	X.XX	X.XX	December	X.XX	X.XX

3.2.3 Forecast performance

The forecast error (or demand uncertainty) measures the difference between the forecast and actual demand. While the forecast performance falls outside the scope of this thesis, it is worth noting that it currently has an impact on the performance of PVM and will be briefly explained in this subsection. The forecast error is a measure of uncertainty and drives all responses to uncertainty, such as safety inventory or excess capacity. The error of the forecast is determined for the 2-month ahead forecast, as this is the lead time required for the supply planning department to take any planning actions. Eq. 3.3 provides the computation of the forecast error $\mathrm{E_T}$ in period T, where $\mathrm{F_T}$ and $\mathrm{D_T}$ resemble the forecast and demand, respectively.

$$\mathrm{E_T} = \mathrm{F_T} - \mathrm{D_T} \tag{3.3}$$

From the forecast error, we can compute the mean absolute percentage error (MAPE), which is the most widely used measure for checking forecast accuracy. It comes under percentage errors which are scale independent and can be used for comparing series on different scales. It also is a good measure of forecast error when the underlying forecast has significant seasonality and demand varies considerably from one period to the next.

$$1 - MAPE = 1 - \frac{1}{n}\sum_{t=1}^{n} \frac{|E_T|}{D_T} * 100\% \tag{3.4}$$

Computing the 1 - MAPE of the entire portfolio resulted in a forecast accuracy of XX.X%. Next to the MAPE, it is also interesting to weigh the errors by sales volume, which results in the weighted mean absolute percentage error (WMAPE). It overcomes one of the drawbacks of MAPE, as it gives more importance to high-selling products. The WMAPE is given by eq. 3.5, which led to a value of XX.X%. This means that the forecast accuracy is more accurate for high-selling products.

$$1 - WMAPE = 1 - \frac{\sum_{t=1}^{n}|E_T| * w_i}{\sum_{t=1}^{n} D_T * w_i} * 100\% \tag{3.5}$$

Next to the forecast accuracy, it is also interesting to analyze the bias of a forecast. The bias is the difference between the total forecast and sales. If the forecast over-estimate sales, the bias is considered positive, and vice versa, the bias is considered negative if the forecast underestimates sales. A bias indicates fundamental problems in the forecast process, which may be the result of manual intervention, wrong parameter settings, or an incorrect model. Eq. 3.6 gives the function used for computing the forecast bias.

$$Forecast\ bias\ \% = \sum \mathrm{E_T} \tag{3.6}$$

This equation led us to a mean forecast bias of -XX.XX tons, which implies that the sales tend to be greater than the forecasted number, so sales are underestimated. The corresponding standard deviation, which is also known as the demand variability, was about XX.XX tons. The standard deviation was quite high as high-selling products experience greater fluctuations than low-selling products. A small remark about this forecast analysis is that there was limited data available, as forecast data has only been saved since 2021. Fig. 3.7 visualizes the monthly standard deviation and bias.

Fig. 3.7 Standard deviation and bias in forecast errors (January 2021 - May 2022).

3.2.4 Inventory Analysis

To date, the inventory status (or stock overview) is examined by taking a snapshot of the monthly opening stock. This stock overview is compared with previous months and assessed on days of stock on hand and shelf life. There is no performance measure that tracks the stockout rate, fill rate, or service levels. To set a benchmark for this research and better understand stock performance, we decided to develop a tool to actively monitor stock availability.

At the start of this research, we requested the supply planners to save all available stock data on a daily basis. Consequently, four months of daily stock overview data were available for this research. Although it is a quite limited sample, it still gives some opportunities for interesting insights. The stock overviews and demand forecast serve as input for the VBA tool, which we created to monitor stock availability.

As it is difficult to measure stockouts because of limitations in data acquisition, we assumed that a stockout occurs if the stock on hand at the DC is less than or equal to XX days. We computed the stock on hand by dividing the available stock by the forecasted demand. This input data was connected through VBA coding. The dashboard updates in real-time (by pressing a button) and visualizes the stock performance accordingly. Appendix A.2 provides a snapshot of the Power BI version of the tool, Appendix A.3 gives the manual of monthly report required for the S&OP meetings.

Table 3.4 presents the stockout performance per DC. It expresses the DC performance per product and per sales volume. From this table, it is evident that the NDC faces the most stockouts and that the other DCs perform is comparable. When we show the stockout as part of the sales volume, instead of stockout per stock keeping unit (product), it can be seen that the impact is much higher. This indicates that more stockouts are experienced for the higher-selling products in the portfolio.

Table 3.4 Stockouts per Distribution Center.

	Total	South	North	Center
Stockout/product	XX.X%	XX.X%	XX.X%	XX.X%
Stockout/volume	XX.X%	XX.X%	XX.X%	XX.X%

It is also relevant to get more insights into the excess stock by analyzing the stock on hand. Although obsoletes and dead stock are very limited, it still provided insights into inefficiencies of PVM's inventory management. Table 3.5 lists the days of stock on hand per DC. The days on hand were the greatest at the SDC because it also serves as a consolidation center. The table also makes clear that the stock on hand tends to be lower for high-selling products, as the rates are greater for stock on hand per product. This can be clarified as the MPQ has a higher relative impact on low-selling products than the popular products.

Table 3.5 Days of stock on hand per Distribution Center.

	Total	South	North	Center
Stock on hand/product	XX.X%	XX.X%	XX.X%	XX.X%
Stock on hand/volume	XX.X%	XX.X%	XX.X%	XX.X%

The trendline in Fig. 3.8 also visualizes increasing stock-on-hand rates for the product with lower sales volumes, as the x-axis is sorted in decreasing order.

Fig. 3.8 Days of stock on hand per core product, sorted by sales volume (decreasing order).

3.2.5 Data Cleaning and Validation

Data entry and acquisition are intrinsically prone to errors both simple and complex. For data sets, the logical solution to this problem is to clean the data in some way. Data cleaning deals with identifying and eliminating errors and inconsistencies in data to improve quality.

Cleaning the data was a laborious and time-consuming task, as several source files from different departments were integrated into this research. The source is a crucial factor for the quality of a real-world data set. The source files often contained redundant data and data in different representation formats. For example, the product codes that are used to differentiate products vary per department.

In order to provide access to accurate and consistent data, consolidation of different data representations and elimination of duplicate or missing information was required. Before consolidating the data, it was necessary to group the products into three classes: seasonal promotion, new product development and core products. This grouping was based on the product descriptions and required the support of stakeholders. In the next step, we filtered the files on redundant data and modified the relevant data to the preferred representation by setting rules with Power Query in Excel.

Now we needed to connect the source files and so we created a lookup table with a "mother mapping" column for the core products. This mother map serves as a key, commonly known as a master product code, which enables the linkage between data from various source files. In other words, the mother map is the umbrella code that links the product information of different departments. The mother map is also beneficial in terms of updating product codes, as changes in product appearance (e.g., slightly adjusted packaging materials) lead to an updated product code, which is prone to errors. As a consequence of this update procedure, multiple product codes resembling the same product are being used simultaneously. But the overarching mother map remains unchanged and robust. Fig. 3.9 depicts the connection of several data sources by mother mapping. Appendix A.4 provides a more detailed example.

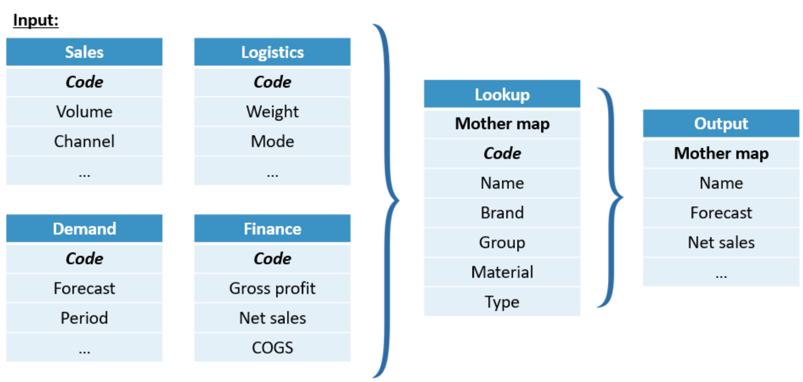

Fig. 3.9 Data table lookup process by mother mapping.

It was essential to validate the final data as the processes associated with the cleaning itself are also error-prone. Therefore, all steps taken to process the data, from the raw version to the final version, were evaluated by stakeholders from the departments concerned. In fact, consulting stakeholders for validation was one of the primary validation techniques in this thesis, i.e., this was critical for the optimization tool created and the conclusions drafted later in this thesis.

3.3 Conclusion of the Context Analysis

This section examined the activities concerned with and the performance of the stock replenishment at PVM. The purpose was (i) to get a deeper understanding of the processes, stakeholders and portfolio related to the replenishment policy, (ii) to fine-tune the research scope (iii) and find research opportunities for improvement.

PVM has three DCs that are monitored independently but offer lateral transshipment between the SDC and the other DCs because of the MPQ's effect on excess production. The SDC serves as both a storage center for southern demand and a consolidation center for northern and central regions. The NDC and CDC only keep enough inventory to fulfill their respective demand.

The stock levels of the DCs are managed by a periodic review policy with a variable lot size, also known as an (R, S) replenishment policy. The review period R and order-up-to-level S are XX and XX days, respectively. The product portfolio consists of three different product types, of which this thesis focuses on the XX core products. Currently, all products are replenished with the same parameter settings. Product volumes and demand variabilities are not considered, and products are not classified based on these or any criteria. The MPQ is pre-determined by the company and therefore out of scope. Then, PVM distinguishes two types of sales channels, namely TT and MT. MT customers have order priority over TT.

This thesis tries to balance the fill rate and logistics cost performance measures. The variability of the demand has a significant influence on this, as the demand faces a high seasonality. The accuracy of the forecast, which is negatively biased, also has an impact. Consequences are getting out-of-stock and holding excessive inventory. Stockouts occur relatively more for high-selling products, while excess inventory is more common for low-selling products. These results are related to each other and confirm the need for balanced inventory levels.

4 Solution Design

This section covers the solution design of the thesis. Firstly, Section 0 applies product classification to PVM's product portfolio. Building upon this, Section 4.2 delves into the development of a conceptual model that forms the foundation of the optimization model, discussed in detail in Section 4.3. Lastly, Section 4.4 provides a summary of the solution design, which serves as the foundation for the optimization.

Fig. 4.1 presents the structure of the solution approach, from which it is visible that this section focuses on the set-up of the solution and Section 5 addresses the outcomes and evaluation of these.

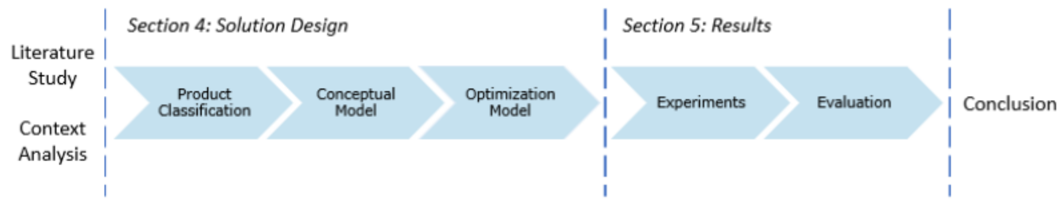

Fig. 4.1 Flowchart of the solution approach.

4.1 Application of Product Classification

As previously outlined in the literature review, the ABC classification provides a means of categorizing products according to their relative value to a business (see Section 2.2). In our first experiment, we utilized this classification to categorize our product portfolio (Section 5.1.1). The combination of both two classifications led to the ABC-XYZ classification, which we examined in our second experiment (Section 5.1.2).

4.1.1 ABC Classification

We computed the annual value per product by eq. 2.1, after which the products were ranked in descending order (from most valuable to least valuable). The products were classified into the three classes, where the products that contributed to 80% of the total annual value were listed in the A-class, B-class products generated 15% of the annual value and C-class products contributed to the remaining 5%.

This led to the classification as presented in Table 4.1. On average, XX.X% of the products were divided into the A-class and XX.X% of the products fell into the B-class. The remaining XX.X% of the portfolio were C-class products. This distribution is similar to the theoretical 20/30/50 power-law. Due to certain products not being sold in the central region, the product portfolio of the CDC was slightly smaller, which is reflected in the distribution presented in the table.

Furthermore, an analysis of the sales distribution across customer channels revealed interesting insights. Among TT sales, A-class products accounted for XX.X% of the sales volume, followed by XX.X% for B-class products, and the remaining XX.X% for C-class products. In contrast, for MT sales, the distribution was more balanced, with XX.X% for A-class, XX.X% for B-class, and XX.X% for C-class products. These findings emphasize the significant contribution of A-class products to TT sales, while MT exhibited a more even distribution across different product classes. These findings underscore the importance of understanding the distinct sales patterns observed across customer channels and the consequential implications for prioritizing and classifying products.

Table 4.1 ABC classification.

	Number of products			
Class	South/North DC	Central DC	Average	Theory
A	XX.X%	XX.X%	XX.X%	20%
B	XX.X%	XX.X%	XX.X%	30%
C	XX.X%	XX.X%	XX.X%	50%
Total	*100%*	*100%*	*100%*	*100%*

4.1.2 XYZ Classification

We used eq. 2.2 to calculate the CV per product based on the monthly demand rates. In line with the theory, we classified products with a CV lower than 0.5 as X-class, with a CV between 0.5 and 1 into the B-class, and the products with a CV greater than 1 into the C-class. Table 4.2 gives the corresponding XYZ classification. As many as XX.X% of the products is X-class, XX.X% of the products belongs to the Y-class and XX.X% is Z-class. From this table it can be concluded that the X-class (low demand variability) and especially the Y-class are extremely large, while the Z-class (high demand variability) is small compared to the theoretical distribution. This means that the portfolio is relatively predictable.

Table 4.2 XYZ classification.

	Number of products			
Class	South/North DC	Central DC	Average	Theory
X	XX.X%	XX.X%	XX.X%	0-10%
Y	XX.X%	XX.X%	XX.X%	10-25%
Z	XX.X%	XX.X%	XX.X%	25+%
Total	*100%*	*100%*	*100%*	-

4.1.3 ABC-XYZ Classification

When we combined the two classifications into the ABC-XYZ classification, we have Tables 4.3 and 4.4 for the SDC, NDC and CDC. These tables were then used to create Table 4.5, which has been visualized in Fig. 4.2.

The tables reveal that the variability of A-class products is the lowest, followed by B-class products. C-class products experience the highest variability and are the only group with Z-class products present. This means that C-class products are more difficult to predict and manage than A- or B-class products.

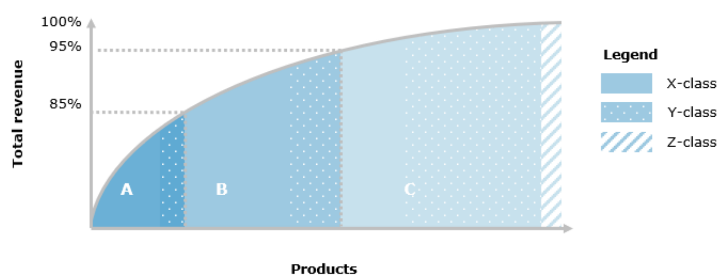

Fig. 4.2 ABC-XYZ classification.

Table 4.3 ABC-XYZ classification South/North DC.

	X	Y	Z
A	XX.X%	XX.X%	XX.X%
B	XX.X%	XX.X%	XX.X%
C	XX.X%	XX.X%	XX.X%

Table 4.4 ABC-XYZ classification Central DC.

	X	Y	Z
A	XX.X%	XX.X%	XX.X%
B	XX.X%	XX.X%	XX.X%
C	XX.X%	XX.X%	XX.X%

Table 4.5 Average ABC-XYZ classification.

Average	X	Y	Z
A	XX.X%	XX.X%	XX.X%
B	XX.X%	XX.X%	XX.X%
C	XX.X%	XX.X%	XX.X%

4.2 Conceptual Model

In this subsection, we present the essential elements of our simulation model through a conceptual model. A conceptual model, a non-technical illustration of the simulation model, offers a simple and comprehensible portrayal of the simulation, while accomplishing the overall objectives of the simulation study (Robinson, 2004). It encompasses six crucial components: objective, inputs (decision variables), outputs, content, assumptions, and simplifications. Fig. 4.3 showcases the components chosen for our conceptual model.

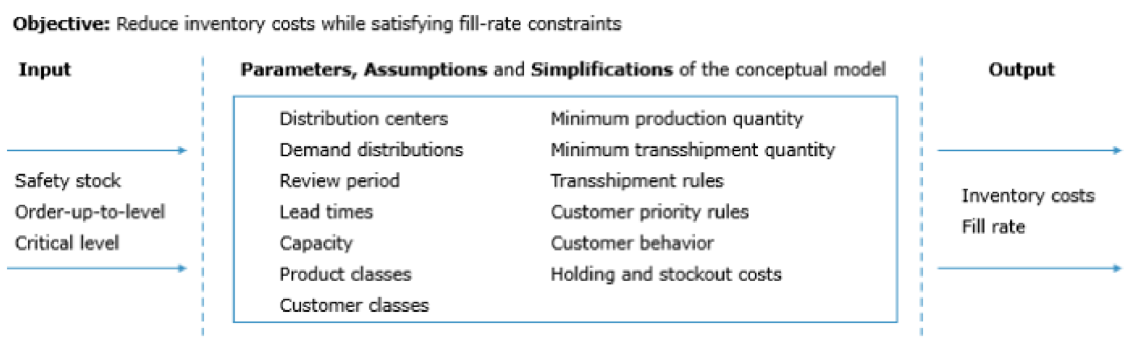

Fig. 4.3 Overview of the key components of the model.

4.2.1 Model Objective

The purpose of this simulation study was to reduce the overall costs associated with inventory management, including inventory and holding cost while satisfying the fill rate constraints. Through an analytical approach, the initial decision variable settings were determined. Then, the simulation model employed a greedy algorithm to search the optimal order-up-to-level S and critical level k settings.

4.2.2 Model Inputs

The simulation model tried to optimize the S and k. We wanted to find the optimal S, specified in days on hand (DOH), by optimizing the safety stock. S is calculated as the sum of the demand during the review period R and lead time L (indicated by $\hat{\mu}_{R+L}$) plus the safety stock, making the safety tock the optimizable element of S. The critical level k, which lies between 0 and S DOH, sets the threshold below which only orders from MT (priority class) are fulfilled, with orders from TT (non-priority) not being satisfied.

In Section 4.1, we divided the product portfolio into 3 classes for the ABC classification experiment. The simulation created a class-based policy for each of the 3 DCs, where each class of the product portfolio is assigned a k and S, with the latter determined by the safety stock. The (R, S, k) policy settings for each class was formed by combining these two variables with the predetermined R. In total, this resulted in 2 decision variables x 3 classes x 3 DCs = 18 decision variables. For the experiments that have an ABC-XYZ classification, 9 classes were used, leading to 2 decision variables x 9 classes x 3 DCs = 54 decision variables.

4.2.3 Model Outputs

We seek to improve the aggregated fill rate β_d of each Distribution Center d and to reduce the inventory costs of the entire system (see Section 3.2). If we were to decentralize costs, the advantages of the transshipment strategy would be lost as each DC, particularly the SDC in this case, would focus on minimizing its own costs rather than the those of the total system. This could potentially lead to higher overall expenses.

The total inventory costs consists of two elements, namely the holding cost and stockout cost. The holding cost encompasses the storage charges per pallet at the third-party DC and includes XX,XXX Vietnamese Dong (VND) per simulation day. To calculate the holding cost h_i for each product i, the storage cost per pallet is divided by the number of products that can be stored on it. t is important to note that in reality, fractionalized pallet storage is not feasible, and the full pallet storage cost is always incurred. In reality, fractionalized pallet storage is not feasible and will always result into full pallet storage cost. However, this constraint is not taken into account as the company prioritizes delivering fractionally filled pallets to customers first, minimizing these expenses. This leads us to a total holding cost of $\sum h_i IL^+_{i,d,t}$, where $IL^+_{i,d,t}$ represents the physical holding inventory of product i at DC d on simulation day t.

To exemplify, consider the scenario outlined in Table 4.6 for a specific product with a holding cots of VND 10,000 per day. The inventory level for this product declines from 4.5 pallets on day 1 to 0 pallets on day 4 due to the demand rate. Calculating the holding cost for the entire period, we have 10,000 $*$ (4.5 + 3.5 + 1.5 + 0.0), which is VND 95,000.

Table 4.6 Stock overview and demand (both in pallets) per simulation day.

Time	Inventory Level (IL^+)	Stockout quantity (IL^-)	Demand
Day 1	4.5	0.0	-
Day 2	3.5	0.0	1
Day 3	1.5	0.0	2
Day 4	0.0	0.1	1.6

The stockout cost $\sum s_i IL^-_{i,d,t}$ is charged proportionally to the shortage level. $IL^-_{i,d,t}$ is the unfulfilled customer order quantity for product i at DC d and time t, while s_i denotes the shortage cost per product i short, which is determined by the stockout quantity multiplied by the difference between revenue and costs of goods sold. The lost sales cost differs per product depending on the product's profitability. The total inventory costs is computed as the sum of the holding cots and stockout cost, given by:

$$\sum (h_i IL^+_{i,d,t} + s_i IL^-_{i,d,t}) \tag{4.1}$$

Again, consider the example from Table 4.6 and assume a shortage cost of VND 1,400,000 is incurred per pallet out-of-stock. This results in a total stockout cost of 1,400,000 * (0 + 0 + 0 + 0.1), which amounts to VND 140,000. Combining this with the holding cost, the total inventory costs equals VND 235,000 (95,000 VND + 140,000 VND).

The aggregated fill rate per DC is computed using eq. 3.1, which can be rewritten as:

$$\beta_d = \frac{\sum_{t\in T, i\in I}(\mu_{i,d,t} - IL^{-}_{i,d,t})}{\sum_{t\in T, i\in I}\mu_{i,d,t}} \tag{4.2}$$

Here, $\mu_{i,d,t}$ represents the demand of product i at DC d and time t. It is imperative to review the fill rate on a regional level, to equally satisfy customer demand nationwide. Additionally, we analyze the cost and fill rate per product class to improve decision-making. In the aforementioned example, the aggregate fill rate would be (4.6 – 0.1) / 4.6, resulting in 97.8%.

4.2.4 Model Content

This subsection elaborates on the content of the conceptual model (see Section 4.2) from two perspectives, namely the scope and level of detail that was included. Robinson (2004) suggests the 80:20 power law as a guideline for determining the content to be included.

The scope covers the core processes in the replenishment policy. We describe these processes through the explanation of the pseudo-code in Fig. 4.4.

```
procedure simulate inventory policy (S, k)
    initialize input ()
    for t ∈ T do
        update starting inventory (t)
        for d ∈ D do
            for i ∈ I do
                modern trade (d, i, t)
                traditional trade (d, i, t, k)
            end for
        end for
        for d ∈ D do
            for i ∈ I do
                substitute traditional trade (d, i, t, k)
            end for
            for i ∈ I do
                if t mod 7 = 0 then
                    shipment quantity (d, i, t, S)
                    order quantity (d, i, t, S)
                else if (t − t_trans.) mod 7 = 0 then
                    shipment quantity (d, i, t, S)
                end if
                update order up to level (d, i, S)
                update critical level (d, i, S, k)
            end for
        end for
        update final inventory (t)
    end for
    compute total costs ()
    compute total fill rates ()
end procedure
```

Fig. 4.4 Pseudo-code of the simulation model.

The model starts by initializing the inputs, which are the stochastic demand, review period and lead times, product classes, customer classes, minimum production and transshipment quantities, and costs components. The decision variables, including the safety stock ss, order-up-to-level S and critical level k, are also determined as part of the initialization stage (line 2).

The model iterates through the simulation days and starts by determining the available inventory (line 4). The MT is served next, updating the inventory and lost sales (line 7). Then, the TT is handled, considering k, and the inventory and stockouts are updated accordingly (line 8). The model checks for substitute products to resolve any stockouts for TT, after which the inventory and stockouts are updated for the last time that day (line 13).

On Thursdays, every DC reviews its inventory levels and place orders to increase it up to the S (lines 16-17). The partial pooling inventory of the SDC is used as far as possible for replenishment (line 18). If this is insufficient, a production order is initiated while ensuring the minimum production quantity is met, potentially resulting in excessive inventory at the SDC.

At the end of each simulation day, S and k are adjusted based on the average demand, and the closing inventory level is calculated (lines 20-21, 24). The next simulation day starts now until all days are simulated. Once the simulation is completed, the stockout and inventory costs, and fill rates are computed (lines 26-27).

4.2.5 Model Assumptions and Simplifications

The focus of this section is on the assumption and limitations made. Assumptions are made in the presence of uncertainties or beliefs about the real world, and simplifications are added to streamline the model development and speed.

Assumptions

(i) The demand (firstly mentioned in Section 3.2) follows a normal distribution pattern, commonly used for high-demand products like candies. To obtain the normally generated demand data, the 3-year historical demand data is adjusted to remove any seasonal and trending influences. This results in deseasonalized and detrended data, from which the mean and standard deviation are calculated and used as inputs for the normal distribution demand. Lastly, the normally distributed demand is multiplied by the relevant seasonal factor to obtain the final demand.

(ii) The forecast error is calculated based on 3-year historical demand and sales data.

(iii) If the production order is less than the minimum production quantity, then the remaining fraction is stored at the SDC and used as partial pooling inventory.

(iv) Each DC places individual orders for required products.

(v) The review period is XX days.

(vi) There are two customer types that place orders daily, namely MT and TT. Orders for MT customers are fulfilled first, followed by orders from TT customers.

(vii) In the event of a stockout, orders can still be partially fulfilled.

(viii) The storage cost is fixed and charged per pallet each day. The cost per product is determined by determining the fraction of the pallet taken up by the specific product.

(ix) The transportation costs are out-of-scope.

(x) For TT, all products that are out-of-stock are replaced in case there is a substitute product from the same product group available. If there are no substitutes available, then a lost sale is assumed or a random alternative substitute, both with a XX% chance. For MT, all products that face stockouts are lost sales. Penalty costs are negligent, for convenience, as their influence is minimal.

Simplifications

(i) The demand has daily interarrival times.
(ii) The production lead time is assumed to be constant at XX days.
(iii) The production capacity is ample, as we are interested in the performance without external factors.
(iv) All transportations are by sea and have constant lead times as specified in Table 3.2.
(v) Emergency shipments are not modeled, as the emphasis is on finding the ideal inventory settings without them. Additionally, the modeling of emergency shipments is complex because they can be triggered both proactively or reactively, and their activation is influenced by case-specific elements such as promotions or third-party agreements.
(vi) In reality, they are triggered using real-time information (sometimes proactive and sometimes reactive), also cases-specific elements are considered to fully utilize the emergency shipments, e.g., by integrated promotions, or third-party agreements, which makes it very complex to model them.
(vii) To trigger a transshipment, the container with a volume of XX cubic meters must be at least 80% full. If this condition is not met, the transshipment is canceled.
(viii) Customer behavior is ignored, such as customers leaving or being dissatisfied. Similarly, the influence of stockouts on market share is ignored for simplicity.

4.3 Optimization Model

In this subsection, we outline the steps involved in initializing the simulation, covering the warm-up period, run length and analytical approximations used to obtain the initial input variables. Furthermore, we explain the optimization algorithm developed to optimize these variables, and the verification and validation methods applied to ensure a reliable model.

4.3.1 Simulation Model

The simulation model was developed using the Python programming language, and the input data was retrieved from an Excel template designed for this thesis.

We arbitrarily set the initial inventory levels at each DC to their average real-life stock levels. This prevents empty starting inventory levels, ensuring that the simulation model reaches steady-state faster. The simulation is then warmed up to ensure that system performance is only based on measurements taken after the system has reached steady-state. Results at the start of the simulation were influenced by the starting conditions, e.g., no pipeline inventory, so the data obtained during this period was not suitable for analysis.

In this thesis, we determined the warm-up period for the simulation by applying Welch's graphical procedure, a well-known technique for handling initialization bias (Law, 2014). After applying this procedure to various system configurations, we identified a warm-up period of 135 weeks, as depicted in Fig. 4.5.

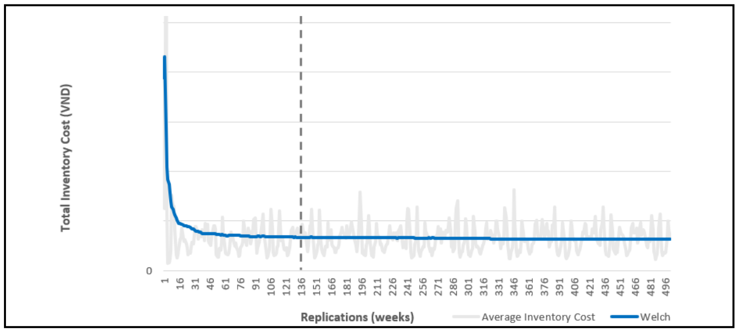

Fig. 4.5 Determining warm-up period.

According to Robinson's (2004) rule of thumb, the run length should be at least 10 times greater than the warm-up period. This would result into a run length of 1485 simulation weeks, which is 11 times greater than 135. But evaluating this run length with the confidence interval test (see eq. 4.3) showed us that a run length of 120 weeks per run is sufficient, which significantly saves simulation time.

To obtain our simulation output, we used the batch means method, which involves running one long simulation run, consisting of the warm-up period plus n runs.

Compared to running multiple independent simulation runs, the batch mean approach is more efficient, as it only requires one warm-up period and provides a better estimate of the system's output with less impact from the initial conditions. But it is crucial for the number of runs, or batches, to be sufficiently large; otherwise, correlations will lead to incorrect confidence intervals. We determined n through the approximate $\mathbf{100(1-\alpha)}$ confidence interval of the simulation output $\bar{X}(n)$, given by (Law, 2014):

$$CI_{95\%} = \bar{X}(n) \pm z_{1-\alpha/2}\sqrt{\frac{\sigma^2(n)}{n}}, \tag{4.3}$$

where $\sigma^2(n)$ resembles the standard deviation of the output from the number of replications per run and α is 5%. This resulted in a total of 4 simulation runs.

4.3.2 Analytical Approximation

For the analytical approach, required for the initial simulation input, the impact of stockouts at both the NDC and CDC on the SDC was taken into consideration when determining S. This was because stockouts at either of these DCs can result in an increase in demand at the SDC due to transshipments. Axsäter and his colleagues (2013) proposed a straightforward and effective method for determining the overflow at the SDC (denoted as d = 1) by computing the mean μ and standard deviation σ of expected lost sales at the other Distribution Centers. Similar to the approach of Kranenburg and Van Houtum (2014), we summed this overflow with the SDC's own demand and defined $\tilde{\mu}_{i,d=1}$ as the annual demand rate for product i at the SDC, i.e.,

$$\tilde{\mu}_{i,d=1} = \left(\mu_{i,1} + \sum\nolimits_{d\in D, d>1} (1-\beta_{i,d})\mu_{i,d}\right) \ \forall i \tag{4.4}$$

The exact analysis of the variance is complicated because $\tilde{\sigma}^2_{d=1}$ depends on the frequency and length of the stockouts at each DC, so we use the approximation:

$$\tilde{\sigma}^2_{i,d=1} = \sigma^2_{i,d=1} + \sum\nolimits_{d\in D, d>1} (1-\beta_{i,d})\sigma^2_{i,d} \quad \forall i \tag{4.5}$$

Based on this information, we determined the mean and standard deviation per product class j. To find the total mean demand, we simply added their individual mean demands:

$$\hat{\mu}_{j,d} = \sum\nolimits_{i\in I} \mu_{i,j,d} \quad \forall j, d \tag{4.6}$$

and for the total standard deviation of the group, we used the formula for the standard deviation of a sum of independent random variables:

$$\sigma_{j,d} = \sqrt{\sum\nolimits_{i\in I} \sigma^2_{i,j,d}} \quad \forall j, d \tag{4.7}$$

Now we determined the initial S using an analytical approach. We used straightforward computations from Section 2.4, where the safety stock equals:

$$ss_{j,d} = \sigma_{j,d,R+L} k_{ss} \quad \forall j, d \tag{4.8}$$

and $\sigma_{j,d,R+L}$ presents the standard deviation of demand under $R+L$, and k_{ss} denotes the safety factor. From this equation S, expressed in DOH, was determined using:

$$S_{j,d} = \frac{\hat{\mu}_{j,d,R+L} + ss_{j,d}}{\hat{\mu}_{j,d,1}} \quad \forall j, d, \tag{4.9}$$

with $\hat{\mu}_{j,d,R+L}$ as demand under the R and L, and $\hat{\mu}_{j,d,1}$ equal to the mean daily demand during $R+L$.

Due to the lack of straightforward analytical approaches to determine k, we decided to start the optimization with no k as the starting point.

4.3.3 Greedy Algorithm

In the first stage of the optimization (line 2 of Fig. 4.6), the algorithm took the analytical solution as an input. But due to the tendency of the analytical solution to overestimate the greedy solution, we made a small modification before using it as an input. The following approach was taken: $\mathrm{Max}(XX, S-\sqrt{S})$. This involved choosing the higher value between XX DOH (as the minimum $L+R$ equals XX) and the analytically determined order-up-to-level S subtracted by its square root. This input was then adjusted in a greedy manner by increased S for the worst performing product class across all DCs .

In stage 2, the worst performing class was found using the equations in lines 4 and 5 of the pseudo code, which divided the actual fill rate by the target fill rate. The algorithm incrementally increased the value of S by one day for this class (line 7), updated the relevant values (line 8-10), and repeated this process across the classes until all fill rate constraints were satisfied.

Stage 3 (line 12) introduced k to the policy optimization by increasing them by 5% for every class and DC one-by-one (as seen in line 16). The algorithm repeatedly checked if the values met the fill rate constraints (lines 19 and 23) and if costs improved (line 23). When these conditions were met, S was decreased by 1 day (line 20). The algorithm then checked if k could be increased by 5% while still satisfying the fill rate constraints and reducing costs (line 32). If not, the iterative optimization was applied to the next product class. Finally, the optimization terminated with line 41, providing the improved performance measures and parameter settings.

```
1.  procedure greedy optimization ()

        Stage 1:
2.      initialize solution space analytically ()

        Stage 2:
3.      β, TC = call simulate inventory policy (S, k)
4.      η = β_{d,j} / β_{target}  ∀d, j
5.      index = arg min η
6.      while β_{d,j=index} < β_{target}  ∀d, j do
7.          S = S + 1
8.          β, TC = call simulate inventory policy (S, k)
9.          η = β_{d,j} / β_{target}  ∀d, j
10.         index = arg min η
11.     end while

        Stage 3:
12.     C_{k,best}, C_{k,new}, C_{S,best}, C_{S,new} = TC
13.     for d ∈ D do
14.         for j ∈ J do
15.             while True do
16.                 k = k + 0.05
17.                 β, TC = call simulate inventory policy (S, k)
18.                 C_{k,new}, C_{S,best} = TC
19.                 while β_{d,j,c} ≥ β_{target}  ∀d, j, c do
20.                     S = S − 1
21.                     β, TC = call simulate inventory policy (S, k)
22.                     C_{S,new} = TC
23.                     if β_{d,j,c} < β_{target}  ∀d, j, c or C_{S,new} > C_{S,best} then
24.                         C_{S,best} = C_{S,new}
25.                     else
26.                         S = S + 1
27.                         β, TC = call simulate inventory policy (S, k)
28.                         C_{k,new} = TC
29.                         break
30.                     end if
31.                 end while
32.                 if β_{d,j,c} < β_{target}  ∀d, j, c or C_{k,new} > C_{k,best} do
33.                     C_{k,best} = C_{k,new}
34.                 else
35.                     k = k − 0.05
36.                     break
37.                 end if
38.             end while
39.         end for
40.     end for
41.     return β*, TC*, S*, k*
42. end procedure
```

Fig. 4.6 Pseudo code of the greedy 3-stage optimization.

4.3.4 Verification and Validation of Simulation

The verification and validation processes ensure the accuracy of the simulation model and results. Verification examines if the model representation of the real system meets the agreed specifications and assumptions, and validation evaluates the accuracy of the model's representation of the real world by ensuring that its implementation matches the conceptualization (Law, 2014; Robinson, 2004).

We used tracing to verify whether the model and greedy algorithm met all requirements, as it is considered the most powerful tool for verification in discrete event simulation (Law, 2014). Tracing involves evaluating the state of the simulation system after each event and comparing whether the outcomes correspond to predictions to ensure that the model works as intended. Another technique we adopted was small problem instance verification, which involved significantly simplifying the model and comparing the results with spreadsheet computations. For our model, we took 3 products, two product groups and classes, and 5 simulation days. Finally, we performed extreme value tests to guarantee logical output values for extreme input values. For instance, we assessed that zero demand led to zero inventory and that removing the minimum production quantity resulted in zero transshipments and a maximum inventory of S at the SDC.

For the model validation, we wanted to replicate real-world inventory management procedures and output data that accurately represented reality. To check this, we inputted the current control policies and accompanying parameters, ran the simulation and compared the performance (see Section 4.4). We then validated various aspects of the model behavior and results with the help of logistical stakeholders and programming experts, through a walk-through. From this, we concluded that the simulation results were consistent with reality. Although the model overestimated the real-world performance as the reality also experiences supply constraints, e.g., insufficient production capacity. We also reviewed the credibility of the model via sensitivity analysis (see section 4.5), and the algorithm's performance using small problem instances (see section 4.6).

4.4 Summary of the Solution Design

In the solution design, we first applied the product classification methods from the literature review to PVM's core product portfolio. As a result, an ABC-classification was implemented, where XX.X% of products were classified as A-class, XX.X% as B-class, and XX.X% as C-class. This distribution closely aligns with the theoretical 20/30/50 power-law distribution. Regarding the XYZ-classification, XX.X% and XX.X% of the products were in X and Y classes, which indicated that the portfolio is relatively predictable. Furthermore, the ABC-XYZ analysis revealed that A-class and B-class products have lower variability compared to the C-class products, with the C-class products displaying the highest variability. Interestingly, a small percentage (XX.X%) of Z-class products were found exclusively within the C-class group.

We then developed the conceptual model that represent the real-world system. The model aimed to optimize the inventory costs, which consists of the holding and holding cost, while satisfying the fill rate constraints. The safety stocks, order-up-to-levels and critical levels were the decision variables of the model. Furthermore, the model incorporated several parameters such as demand and lead times, and various assumptions and simplifications were made for modeling purposes

The simulation model was developed using the Python programming language, and the input data was retrieved from an Excel template designed for this thesis. To handle the initialization bias, we used a warm-up period of 135 weeks. We generated our output using one long simulation run of the warm-up period plus 4 batches of 120 weeks.

Fig. 4.7 illustrates the 3-stage optimization approach. In stage 1, the initial order-up-to-levels were determined using analytical approximations. These were later refined through simulation using a greedy algorithm in stages 2 and 3. In the second stage, the algorithm attempted to identify the optimal order-up-to-levels by incrementally increasing them until target fill rates were achieved and the inventory costs and costs could no longer be reduced. The third stage introduced the critical level policy, which involved gradually increasing the parameter k. This adjustment allowed for potential reductions in the parameter S, as long as the target fill rates were maintained and costs did not increase. The algorithm iterated through each product class, moving on to the next when it was no longer possible to further increase k or reduce S. The optimization process terminated upon completion of the final product class.

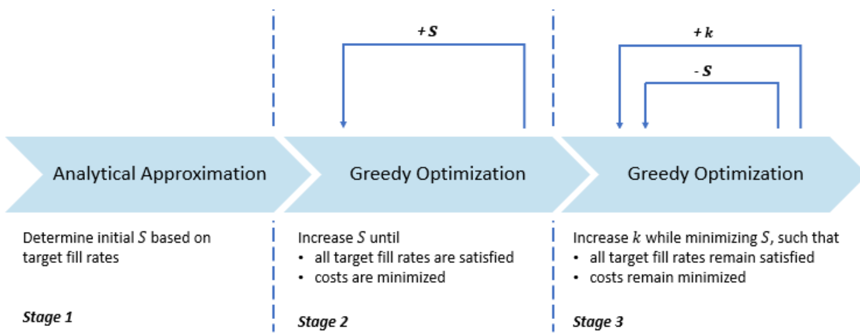

Fig. 4.7 Flow of the 3-stage optimization approach.

5 Solution Execution

This part of the solution elaborates on the results in Section 5.1 and proceeds to evaluate them in Section 5.2. The performance of the greedy algorithm is reviewed in Section 5.3, and the key findings of this section are summarized in Section 5.4.

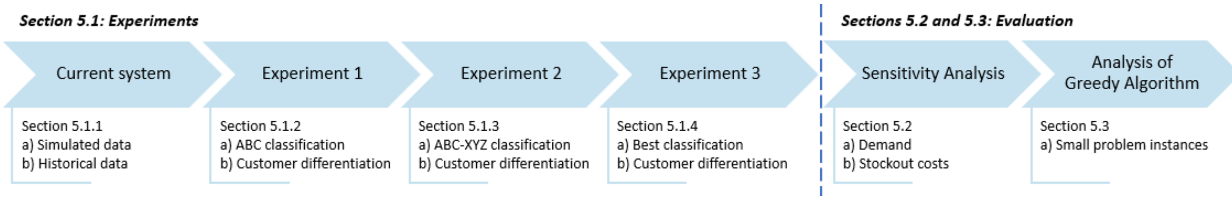

Fig. 5.1 Outline of the solution execution.

Fig. 5.1 provides more detail on the structure of this section. From this figure it is visible that Section 5.1 focuses on the results and Sections 5.2 and 5.3 cover the evaluation of the solution.

5.1 Results of the Optimization

Before conducting experiments, we analyzed the simulation under its current settings.

5.1.1 Current System Configuration

To evaluate the performance of the simulation model, it was key to run it with the current real-world settings. We ensured this by setting the general order-up-to-level S to XX DOH, while excluding the critical level policy and product classification.

The results are in Tables 5.1-5.3, where the fill rates between parentheses imply the performance with product substitutions included. Recall that product substitutions are only possible for the TT customers in case of a stockout (see Section 3.1.5). These substitutions are random with a XX% chance.

Tables 5.1 and 5.2 show that the normally distributed data gives an overall performance of XX.X%, which falls within the real-world performance range of XX.X% to XX.X%. The NDC has the lowest performance in both the simulation and the real-world scenarios. The SDC has a high fill rate of XX.X%, which can be explained because the stock levels tend to exceed the order-up-to-level as a result of the MPQ, i.e., extra buffer inventory. In reality, however, this excess inventory tends to have a smaller effect as proactive transshipments prevent extreme accumulation, leading to a more balanced performance across DCs.

Next, the demand data is of utmost importance in this simulation study. We used a normal distribution to replicate the demand. To assess the normal distribution's ability to mimic real demand, we also used historical data from the past 3 years. The results in Tables 5.1-5.3 reveal that the normally distributed data produced slightly higher fill rates and lower costs compared to the real data. As Robinson (2014) noted, these differences are not surprising when the historical data set is relatively small.

In general, the fill rates follow similar patterns for both datasets. Appendix A.5 gives a graph that visualizes both demand data patterns, from which it is visible that both data sets behave similar but the historical data has a slightly higher total demand. Consequently, this results in higher holding and holding cost, as listed in Table 5.3. Section 5.2.1 gives more details on the demand analysis.

Table 5.1 DC fill rates under normally distributed and historical data.

	Total	SDC	NDC	CDC
Simulated data	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X% (XX.X%)
Historical data	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X% (XX.X%)

Table 5.2 Channel fill rates under normally distributed and historical data.

	Traditional Trade	Modern Trade
Simulated data	XX.X% (XX.X%)	XX.X%
Historical data	XX.X% (XX.X%)	XX.X%

Table 5.3 Costs (in millions) under normally distributed and historical data.

	Inventory costs	Holding cost	Holding cost
Simulated data	VND XX,XXX	VND XX,XXX	VND XX,XXX
Historical data	VND XX,XXX	VND XX,XXX	VND XX,XXX

5.1.2 Experiment 1: ABC Classification

The first experiment focused on the policy settings for the ABC classification.

Stage 1: Analytical Initialization

The optimization process started by using the analytical computations from Section 5.3.1 and the total target fill rates of Table 5.5 to determine the control parameters S. Table 5.4 presents the outcomes in DOH, where ~SDC represents S including the lateral pooling effect of the other DCs. Comparing the ~SDC and SDC shows that the pooling effect has a relatively small impact with a maximum difference of XX DOH. From the table, it is evident that S tend to increase as the sales volumes decrease. This can be attributed to the fact that these classes include the most volatile products, whose demand is less predictable.

Table 5.4 Order-up-to-level (in DOH) per class and Distribution Center by the analytical approach (stage 1).

Class	A	B	C
~SDC	XX	XX	XX
SDC	XX	XX	XX
NDC	XX	XX	XX
CDC	XX	XX	XX

Stages 2 and 3: Greedy Optimization

As explained in Section 4.3.1, to ensure the simulation model has flexibility in determining the parameter settings, we re-initialized S per class by taking $Max(XX, S-\sqrt{S})$ as a modified input. For example, the value of S for class B at the SDC was modified from XX to XX by applying $Max(XX, XX-\sqrt{XX})$. This step prevents the greedy algorithm from overestimating due to the potentially overestimated approximations from the analytical approach. The overestimation occurred because the analytical computations did not consider the MPQ, transshipments, and substitutions, which provide flexibility.

The second stage of the optimization involved finding the optimal S for each class by gradually increasing them until the fill rate constraints were met and inventory costs were no longer reduced.

Table 5.5 displays the target fill rate for the product and customer classes. The third, and final stage, introduced the critical level k, and the algorithm identified opportunities to lower the S for each class.

Table 5.5 Target fill rate per product class.

Total			Traditional Trade			Modern Trade		
A	B	C	A	B	C	A	B	C
98%	97%	94%	97%	96%	93%	99%	98%	95%

Initially, we incorporated the MPQs in the entire optimization process. However, it turned out that S at the SDC was set too low (i.e., even below the expected demand during the lead time and review period) due to the excess production. In other words, the MPQ ensured more than sufficient inventory was available to satisfy the southern demand and achieve the target fill rates. To exemplify, Appendix A.6.provides these optimizations for the ABC-XYZ classification.

For the actual optimization, we pre-determined S for the SDC without considering the MPQ, and then calculated S for the other DCs, including the MPQ. As S was almost always met, the effect of increasing them to the pre-determined values was negligible compared to pre-determining them. But it was relevant to the company in order to set an appropriate active trigger point for production orders at the SDC.

The results of the second stage of optimization are listed in Table 5.6. From this table, it is visible that the values are slightly lower compared to the analytical approximation. Also the pattern is different, as the MPQs ensures sufficient inventory at the SDC, which can also be used for transshipments, S tends to be lower as the sales volumes decrease. Regarding the performance, it is visible from Table 5.9 that although the MT was given priority and higher target fill rates, its performance was found to be lower than TT. This can be attributed to the fact that TT has a significant presence in the A-class, with A-class products accounting for XX.X% of its sales (see Section 4.1). In contrast, MT is more evenly distributed across the different product classes.

The third step optimization, introducing k, did not lead to significant improvement, as visualized in Tables 5.8-5.10. The algorithm only selected critical levels of XX.X% for C-class products at the Northern and Southern DC. But the values of S were never decreased, as doing so would have led to either unsatisfied fill rate constraints or increased costs. Then we also performed a t-paired test, with 95% confidence, to analyze whether there was a performance increase in stage 3 compared to stage 3 (see Appendix A.7). From this test we could not conclude there was. Additionally, the MT fill rate, and S remained unchanged.

Table 5.6 Order-up-to-level (in DOH) per class and Distribution Center by the algorithmic approach (stages 2 and 3).

Class	A	B	C
SDC	XX	XX	XX
NDC	XX	XX	XX
CDC	XX	XX	XX

Table 5.7 Critical levels (as fraction of S) per class and Distribution Center by the algorithmic approach (stage 3).

Class	A	B	C
SDC	XX	XX	XX
NDC	XX	XX	XX
CDC	XX	XX	XX

Table 5.8 DC fill rates under optimization stages 2 and 3.

	Total	SDC	NDC	CDC
Stage 2	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X% (XX.X%)
Stage 3	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X% (XX.X%)

Table 5.9 Channel fill rates under optimization stages 2 and 3.

	Traditional Trade	Modern Trade
Stage 2	XX.X% (XX.X%)	XX.X%
Stage 3	XX.X% (XX.X%)	XX.X%

Table 5.10 Costs (in millions) under optimization stages 2 and 3.

	Inventory costs	Holding cost	Holding cost
Stage 2	VND XX,XXX	VND XX,XXX	VND XX,XXX
Stage 3	VND XX,XXX	VND XX,XXX	VND XX,XXX

5.1.3 Experiment 2: ABC-XYZ Classification

In addition to the ABC classification, we also determined the near-optimal policy settings for the ABC-XYZ classification.

Stage 1: Analytical Initialization

Using the analytical computations and total target fill rates of Table 5.12, we got the initial policy settings as provided in Table 5.11. Generally, the values for S are a bit higher compared to Table 5.4. This could be explained by the fact that the majority of sales comes from the X-class, followed by the Y-class (see Section 4.1.3). As these are now split up, due to the XYZ classification, S decreased. Take the B class of the SDC for example, which is XX DOH for the ABC and XX and XX DOH for BX and BY, respectively.

Table 5.11 Order-up-to-level (in DOH) per class and Distribution Center by the analytical approach (stage 1).

Class	AX	AY	AZ	BX	BY	BZ	CX	CY	CZ
~SDC	XX	XX	XX	XX	XX	XX	XX	XX	XX
SDC	XX	XX	XX	XX	XX	XX	XX	XX	XX
NDC	XX	XX	XX	XX	XX	XX	XX	XX	XX
CDC	XX	XX	XX	XX	XX	XX	XX	XX	XX

Stages 2 and 3: Greedy Optimization

The second stage of the optimization involved finding the optimal S for each class by gradually increasing them until the fill rate constraint were met and inventory costs were no longer reduced. In line with experiment 1, Table 5.13 shows lower values for S compared to stage 1 (see Table 5.11 for stage 1).

Table 5.12 Target fill rate per product class.

	Total			Traditional Trade			Modern Trade		
	X	Y	Z	X	Y	Z	X	Y	Z
A	99%	97%	95%	98%	96%	94%	99%	98%	97%
B	97%	95%	93%	96%	94%	92%	98%	97%	96%
C	95%	93%	91%	94%	92%	90%	97%	96%	95%

The third stage of the optimization resulted in only a minimal increase in performance compared to the second, which can be observed in Tables 5.14-5.16. From Table 5.16, it is visible that the inventory costs increased marginally by XX.X% compared to the ABC classification. The algorithm only selected critical levels between XX.X% and XX.X% of S for the C class. But the values of S were never decreased, as doing so would have led to either unsatisfied fill rate constraints or increased costs.

Table 5.15 demonstrates a slight increase in the MT performance, while the TT channel experienced a small decrease. Although MT was given priority and higher target fill rates, its performance was found to be lower than TT. This can be attributed to the fact that TT has a significant presence in the A-class, with A-class products accounting for XX.X% of its sales (see Section 4.1). In contrast, MT is more evenly distributed across the different product classes.

Table 5.13 Analytically determined order-up-to-level per class and Distribution Center (stages 2 and 3).

Class	AX	AY	AZ	BX	BY	BZ	CX	CY	CZ
SDC	XX	XX	XX	XX	XX	XX	XX	XX	XX
NDC	XX	XX	XX	XX	XX	XX	XX	XX	XX
CDC	XX	XX	XX	XX	XX	XX	XX	XX	XX

Table 5.14 D DC fill rates under optimization stages 2 and 3, including experiment 1.

	Total	SDC	NDC	CDC
Experiment 1	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X% (XX.X%)
Experiment 2 (stage 2)	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X% (XX.X%)
Experiment 2 (stage 3)	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X% (XX.X%)

Table 5.15 Channel fill rates under optimization stages 2 and 3, including experiment 1.

	Traditional Trade	Modern Trade
Experiment 1	XX.X% (XX.X%)	XX.X%
Experiment 2 (stage 2)	XX.X% (XX.X%)	XX.X%
Experiment 2 (stage 3)	XX.X% (XX.X%)	XX.X%

Table 5.16 Costs (in millions) under optimization stages 2 and 3, including experiment 1.

	Inventory costs	Holding cost	Holding cost
Experiment 1	VND XX,XXX	VND XX,XXX	VND XX,XXX
Experiment 2 (stage 2)	VND XX,XXX	VND XX,XXX	VND XX,XXX
Experiment 2 (stage 3)	VND XX,XXX	VND XX,XXX	VND XX,XXX

For the third experiment, we decided to proceed with the ABC classification instead of the ABC-XYZ, as the outcomes are similar, but the ABC classification is easier to implement for PVM.

5.1.4 Experiment 3: Complete Pooling

Our previous experiments involved using PVM's standard partial pooling method, which permitted the transfer of moving inventory from the SDC to the other DCs when its IP exceeded threshold S. In this experiment, we cancelled the threshold and examined the potential of complete pooling. Most striking is that the fill rates are similar to experiments 1 and 2, but the inventory costs increased as extra safety stock was needed to meet the target fill rates at the SDC, primarily for the higher demand classes A and B.

Table 5.17 shows that S increased for the SDC, but remained unchanged for the other DCs when rounding them to DOH. They slightly lowered, but too little to make any impact on the rounded S. The stock requirements at the SDC increased drastically as it now needed to share all of its stock with the other DCs. Consequently, the holding cost increased as well (see Table 5.20). The stockout cost was reduced due to a higher overall fill rate (see Tables 5.18 and 5.19). But it did not decrease enough to compensate for the increased holding cost, which resulted in increased inventory costs compared to experiments 1 and 2.

Table 5.17 Order-up-to-level per class and Distribution Center by the algorithmic approach (phase 2).

Class	A	B	C
SDC	XX	XX	XX
NDC	XX	XX	XX
CDC	XX	XX	XX

Again, introducing the critical level policy did not lead to notable changes, so we will not consider it in this subsection.

Table 5.18 DC fill rates under optimization stage 2, including experiment 1.

	Total	SDC	NDC	CDC
Experiment 1	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X% (XX.X%)
Experiment 3	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X% (XX.X%)

Table 5.19 Channel fill rates under optimization stage 2, including experiment 1.

	Traditional Trade	Modern Trade
Experiment 1	XX.X% (XX.X%)	XX.X%
Experiment 3	XX.X% (XX.X%)	XX.X%

Table 5.20 Costs (in millions) under optimization stage 2, including experiment 1.

	Inventory costs	Holding cost	Holding cost
Experiment 1	VND XX,XXX	VND XX,XXX	VND XX,XXX
Experiment 3	VND XX,XXX	VND XX,XXX	VND XX,XXX

5.2 Sensitivity Analysis on Simulation

The sensitivity analysis identifies the "sensitivity" of the critical input parameters in this thesis, which gives us insights into the robustness and reliability of the simulation model. Section 5.2.1 reviews the demand parameter, while Section 5.2.2 concentrates on the holding cost.

5.2.1 Demand Sensitivity Analysis

In this simulation, the fundamental input was the demand, which followed a normal distribution. Section 4.4 already compared the normally distributed to the historical demand, which gave similar results but a slightly lower performance for the historical demand. A chi-squared test in Appendix A.5 confirmed the accurate representation of the demand.

This subsection further analyzed the influence of the demand data of the simulation outcomes and conducted a sensitivity analysis on the mean and standard deviation. The mean was incrementally increased by 5% from the current mean up to 15%, as future sales are expected to increase.

Table 5.21 gives the outcomes of this analysis. From this table it is visible that improved policy provided the same performance, but the inventory costs slightly increased due to the rising demand. Applying the greedy algorithm to the adjusted data resulted in the same improved policy settings. Therefore, the "Updated improved policy" column only contains blanks.

Table 5.21 Results of sensitivity analysis on mean demand.

	Improved policy		Updated improved policy	
Mean	ill rate	Inventory costs (M)	Fill rate	Inventory costs (M)
Current	*XX.X% (XX.X%)*	*VND XX,XXX*	-	-
+5%	XX.X% (XX.X%)	VND XX,XXX	-	-
+10%	XX.X% (XX.X%)	VND XX,XXX	-	-
+15%	XX.X% (XX.X%)	VND XX,XXX	-	-

Similarly, the standard deviation was adjusted by first decreasing it by -5%, and then increasing it incrementally by 5% until it reached +15%. Table 5.22 indicates that the performance of the improved solution decreased marginally when the standard deviation was adjusted. Slight enhancements were observed when the algorithm was applied to the demand with the adjusted standard deviation.

Table 5.22 Results of sensitivity analysis on standard deviation demand.

	Improved policy		Updated improved policy	
Standard deviation	Fill rate	Inventory costs (M)	Fill rate	Inventory costs (M)
-5 %	XX.X% (XX.X%)	VND XX,XXX	XX.X% (XX.X%)	VND XX,XXX
Current	*XX.X% (XX.X%)*	*VND XX,XXX*	-	-
+5%	XX.X% (XX.X%)	VND XX,XXX	XX.X% (XX.X%)	VND XX,XXX
+10%	XX.X% (XX.X%)	VND XX,XXX	XX.X% (XX.X%)	VND XX,XXX
+15%	XX.X% (XX.X%)	VND XX,XXX	XX.X% (XX.X%)	VND XX,XXX

Finally, tested an extreme scenario to assess the simulation's robustness. The scenario involved increasing both the mean and standard deviation of the demand by 15%. Table 5.23 shows the outcomes of this final test on demand. The improved policy led to a modestly decreased performance under the extreme scenario. Re-applying the greedy algorithm to this scenario resulted in an increased performance, especially with regard to the inventory costs.

Table 5.23 Results of sensitivity analysis on mean and standard deviation demand.

	Improved policy		Updated improved policy	
Scenario	Fill rate	Inventory costs (M)	Fill rate	Inventory costs (M)
Current	*XX.X% (XX.X%)*	*VND XX,XXX*	-	-
Extreme	XX.X% (XX.X%)	VND XX,XXX	XX.X% (XX.X%)	VND XX,XXX

5.2.2 Holding cost Sensitivity Analysis

Table 5.24 summarizes the results of the sensitivity analysis conducted on the holding cost, which was another critical parameter. We varied the holding cost from the current value up to +5% with step sizes of 1.25%. In addition, we included a value of 10% to analyze the impact whenever the holding cost were relatively high. The near-optimal policy provided the best outcomes up to an increased lost sales costs of +2.5%. But beyond this threshold, we found updated improved policy settings that led to a slightly increased performance.

Table 5.24 Results of sensitivity analysis on lost sales costs.

	Improved policy		Updated improved policy	
Stockout cost	Fill rate	Inventory costs (M)	Fill rate	Inventory costs (M)
Current	*XX.X% (XX.X%)*	*VND XX,XXX*	-	-
+1.25%	XX.X% (XX.X%)	VND XX,XXX	-	-
+2.50%	XX.X% (XX.X%)	VND XX,XXX	-	-
+3.75%	XX.X% (XX.X%)	VND XX,XXX	-	VND XX,XXX
+5.00%	XX.X% (XX.X%)	VND XX,XXX	-	VND XX,XXX
+10.0%	XX.X% (XX.X%)	VND XX,XXX	-	VND XX,XXX

5.3 Evaluation of Greedy Algorithm

It is important to evaluate the greedy algorithm's performance to assess the robustness and optimality of the solution, which is crucial for understanding the limitations and strengths of the model. In this evaluation, we focused on the ABC classification as it involves fewer decision variables compared to the ABC-XYZ classification. However, even with the reduces number of variables, achieving the optimal inventory policy within a reasonable time frame remained unattainable. This was primarily due to the large number of decision variables (9 in total, i.e., 3 DCs x 3 classes) and the large number of potential values for each. Assuming every variables has a maximum of 6 potential values, there are 6 to the power of 9 possible policy configurations, resulting in more than 10 million combinations. With a CPU time of about 6.5 seconds, running the simulation 10 million times would take over 50 years.

One approach to evaluate the performance of the simulation is to create small problem instances by reducing the size of the problem instance and obtaining an optimal solution as a benchmark. By simplifying the model, we could optimize it completely, which allowed us to analyze the behavior of the heuristic and compare it to the optimal solution.

Accordingly, we created 5 small problems and by reducing the number of products, product classes, DCs, or a combination of these. This substantially decreased the computational time required to achieve the optimal solution, with new times varying from 10 hours to 24 hours. The optimal solution is to the one that minimizes inventory costs while meeting the fill rate constraints.

We assessed the performance of the greedy algorithm by measuring the discrepancy between its solution and the optimal solution. That is:

$$error\ (\%) = \frac{(TC^{Greedy} - TC^{Best})}{TC^{Greedy}} * 100, \tag{5.1}$$

where TC^{Greedy} are the inventory costs obtained using the method described in Section 4.3.3. On average, the greedy algorithm yielded a solution with an error rate of 7.27% (see Table 5.25). Among the trials conducted, the worst performing trial had an error rate of 8.92%, while the best solution achieved an error rate within 5.73% of the optimal solution. By applying the average solution error to the results obtained from the improved ABC classification (see Section 5.1.2), we identified an unutilized potential of approximately VND XX,XXX M. This unutilized potential represents the potential cost savings that could be achieved if the inventory management approach was optimized further.

When analyzing the CPU times, it is challenging to draw definitive conclusions. However, there is a noticeable trend of performance increasing as the CPU time increases.

Table 5.25 Results of small problem instances analysis on inventory costs.

Problem instance	Optimal inventory costs (in million)	Greedy inventory costs (in million)	Solution error	Decision variables	CPU time
1	VND XX,XXX	VND XX,XXX	5.73%	6	23:57 hrs.
2	VND XX,XXX	VND XX,XXX	7.45%	6	18:28 hrs.
3	VND XX,XXX	VND XX,XXX	8.92%	6	10:05 hrs.
4	VND XX,XXX	VND XX,XXX	6.96%	4	10:32 hrs.
5	VND XX,XXX	VND XX,XXX	7.32%	6	21.45 hrs.
Average	-	-	*7.27%*	-	-

It is important to keep in mind that small instance testing may not be a complete evaluation of the greedy algorithm's performance. As the size and complexity of the problem instance increase, the behavior of the algorithm can vary, resulting in a different relative gap to optimal performance.

Finding research specifically focused on applying a greedy algorithm to dynamic fast-moving multi-item multi-location problems, like the one addressed in this thesis, was challenging. However, we did find some insights into the performance of greedy algorithms in inventory management from other studies. Shu (2010) reported errors within 3% and 4% when applying a greedy algorithm to warehouse-retailer network problems. Drent and Arts (2020) examined the percentage error of the greedy algorithm for a spare parts inventory problem with relaxed constraints and achieved an average performance of 2.4% with a worst performance of 8.4%. Van der Heide and Van Foreest (2017) developed a near-optimal greedy algorithm for minimizing costs in a spare parts base stock policy and achieved performance within 0.2% of optimality with a maximum error of 5.94%.

Comparing these findings to our research, we observe lower solution errors in the studies mentioned. However, it is important to note that we could not find any relevant research that utilized real-world case study data like ours. The utilization of real-world data and the correlations of clustering data into product classes within our thesis may have an impact on the performance of the greedy algorithm.

5.4 Summary of Solution Design

Initially, we ran the simulation model with the current system settings and observed a similar performance to the real world. Next, we conducted three experiments to find the best policy settings, the results of which are provided in Table 5.26. It is important to note that the fill rates in parentheses indicate the performance excluding product substitutions.

Experiment 1 focused on the optimum policy settings for the straightforward ABC classification. Experiment 2 tried to optimize the ABC-XYZ classification. Both experiments resulted in an increased fill rate of +7.7% and reduced inventory costs of about VND XX,XXX M, meaning a -62% inventory costs reduction. The experiments demonstrated that increasing inventory levels, which in turn led to higher holding cost, resulted in a significant decrease in stockout cost. Consequently, this improvement in managing stockouts contributed to the overall enhancement of total inventory costs.

For the third experiment, we decided to proceed the optimization with the optimization ABC classification instead of the ABC-XYZ, as the results were similar, but the former approach is more straightforward for the company. Experiment 3 involved considering the transshipment strategy, or inventory pooling, and tried to identify opportunities for improvement by implementing complete pooling instead of the currently applied partial pooling. However, we noticed a decrease in performance compared to the other experiments, primarily due to a significantly increase in holding cost at the SDC.

Table 5.26 Total costs and fill rates under the current system and experiment settings.

	Inventory costs (in M)	Holding cost (in M)	Holding cost (in M)	Fill rate	Fill rate TT	Fill rate MT
Current	VND XX,XXX	VND XX,XXX	VND XX,XXX	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X%
Experiment 1	VND XX,XXX	VND XX,XXX	VND XX,XXX	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X%
Experiment 2	VND XX,XXX	VND XX,XXX	VND XX,XXX	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X%
Experiment 3	VND XX,XXX	VND XX,XXX	VND XX,XXX	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X%

We validated the simulation using several techniques, including tracing, small problem instance verification, extreme value tests, historical data comparison, stakeholders walk-through, sensitivity analysis. The sensitivity analysis on demand and holding cost showed that the performance remained stable even though the parameters were changed.

To assess the performance of the greedy algorithm, we solved 5 small problem instances since it was not feasible to determine the optimal solution for the entire product portfolio due to computation time constraints. The optimality gap was defined as the percentual difference between the inventory costs of the algorithm's solution and the optimal solution. The results from the small problem instances showed that the greedy algorithm was able to produce solutions with errors ranging from 5.73% to 8.92%, with an average error rate of 7.27%. A comparison to greedy algorithms adopted in other research in inventory management showed that this error was relatively high. One possible explanation for this discrepancy is that this thesis utilizes different system dynamics, specifically through the use of discrete event simulation, and incorporates real-world data.

6 Conclusion

Perfetti Van Melle's (PVM) inventory management approach has resulted in excess inventory and stockouts due to its reliance on practical experiences instead of established supply chain management theory. To address this issue, this thesis aimed to determine near-optimal replenish policy settings that minimize inventory costs while maximizing product availability.

The company manages its stock levels at its distribution centers (DCs) using an (R, S) replenishment policy, in which the inventory position is raised to the order-up-to-level S every XX review moment R. The order-up-to-level is the product of the safety stock plus the demand during review period and lead time. In the current situation, PVM adopts a standardized approach with an order-up-to-level of XX days of inventory on hand (DOH) for each product and DC. However, this approach does not take into account specific characteristics of individual products or DCs, indicating the lack of a tailor-made policy.

To improve PVM's replenishment policy, we conducted a literature study and context analysis, identifying product and customer classifications as opportunities for improvement. As such, we applied the ABC, XYZ and ABC-XYZ classifications to PVM's core product portfolio. PVM has already segmented its core customer into Traditional Trade (TT) and Modern Trade (MT) classes. Our literature review revealed that introducing a critical level policy could be beneficial in serving different customer classes. This led to an (R, S, k) inventory policy, where low-priority TT customers are not served once the inventory falls below the critical level k. In general, the critical level policy results in a lower overall fill rate but higher fill rates for the priority customers. Furthermore, we aimed to evaluate the effectiveness of PVM's existing transshipments strategy, known as a partial pooling policy, and investigate whether adopting a complete pooling strategy would yield better performance. In the current partial pooling approach, only excess inventory is used to facilitate faster shipments from the SDC to the other DCs. In contrast, a complete policy would utilize the entire stock for transshipments.

We established a three-stage optimization approach to determine near-optimal policy settings for the (R, S, k) policy. In the first stage, we used analytical approximations to compute the initial control parameters. Secondly, we optimized S by incrementally increasing the safety stock using a greedy algorithm. The final stage introduced the critical levels and searched for opportunities to lower the order-up-to-levels while improving MT performance.

Next, we conducted three experiments to test the effectiveness of the methods. The first experiment focused on the replenishment policy using the ABC classification, while the second experiment expanded it to incorporate the ABC-XYZ classification. The results of these experiments, including costs, fill rates, and fill rates with substitutions (indicated in parentheses), are presented in Table 6.1. It is important to note that the possibility of substitutions applies only to TT customers. When demand exceeds the available inventory, the customer facing an out-of-stock situation chooses a substitute product if an alternative from the same type is available

From Table 6.1, it is apparent that the performance of the ABC and ABC-XYZ classifications was almost identical. This similarity can be explained due to the dominance of the A and B classes, representing 95% of total sales and having relatively low demand fluctuations. As the Z-class was only present in the C class, which accounted for only a small fraction of total sales, the impact of the XYZ classification was reduced. Hence, for the third experiment, we exclusively considered the ABC classification because of its ease of implementation for PVM.

Introducing the critical level policy did not lead to notable improvements either. The stringent target fill rates often could not be met. Additionally, inventory costs increased due to more stockouts at the TT channel, which prevented a reduction of the order-up-to-levels.

In the third experiment, we explored the possibility to utilize a complete pooling strategy as opposed to the current partial pooling approach. However, implementing complete pooling resulted in higher inventory costs due to the significant increase in holding cost at the Southern DC. Despite slightly improved fill rates, and thus reduced stockouts, no substantial performance enhancements were observed.

Table 6.1 Total costs and fill rates under the current system and experiment settings.

	Inventory costs (in M)	Holding cost (in M)	Holding cost (in M)	Fill rate	Fill rate TT	Fill rate MT
Current	VND XX,XXX	VND XX,XXX	VND XX,XXX	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X%
Experiment 1	VND XX,XXX	VND XX,XXX	VND XX,XXX	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X%
Experiment 2	VND XX,XXX	VND XX,XXX	VND XX,XXX	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X%
Experiment 3	VND XX,XXX	VND XX,XXX	VND XX,XXX	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X%

In summary, we have determined that a partial pooling policy without critical levels is the best choice for PVM based on the greedy algorithm. Additionally, PVM should adopt the ABC classification in its inventory management strategy with the order-up-to-levels, as shown in Table 6.2.

Table 6.2 Order-up-to-levels (in DOH) per class and distribution center.

Class	A	B	C
Southern DC	XX	XX	XX
Northern DC	XX	XX	XX
Central DC	XX	XX	XX

These policy settings led to the fill rates and costs as provided in Tables 6.3 and 6.4, respectively.

In Table 6.3, the total fill rate increased by 7.7%, the TT fill rate by 7.8% and the modern trade fill rate by 6.9%. Modern trade had the lowest performance despite having higher target fill rates and priority. This can be explained as more TT products are mostly present in the A-class, which is subject to even higher fill rates.

Table 6.4 gives more insights into the trade-off made between holding inventory and facing stockouts. It revealed that an increased holding cost, as a result of more inventory on hand, lowered the stockout cost significantly. Consequently, the total inventory costs were reduced by VND XX,XXX M annually, equivalent to 62%. Assuming the state without external disruptions, such as supply constraints, is realized XX% of the time, this leads to a total reduction of VND X,XXX M.

Table 6.3 Fill rates of the near-optimal solution.

Policy	Total	Traditional Trade	Modern Trade
Current	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X%
Improved	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X%

Table 6.4 Costs (in millions) of the near-optimal solution during the unconstrained state.

Policy	Inventory costs	Holding cost	Stockout cost
Current	VND XX,XXX	VND XX,XXX	VND XX,XXX
Improved	VND XX,XXX	VND XX,XXX	VND XX,XXX
Difference	- VND XX,XXX	+ VND XX,XXX	- VND XX,XXX

We validated the simulation using several techniques, including tracing, small problem instance verification, extreme value tests, historical data comparison, stakeholders walk-through, and sensitivity analysis. The sensitivity analysis demonstrated the robustness of the simulation, as altering the key input parameters (i.e., the demand and holding cost) did not remarkably impact the performance. From the small problem instances, we conclude that the solution had a relatively high optimality gap of 7.27%.

6.1 Recommendations

Our main advice is for PVM to implement the policy settings from Table 6.2 in SAP to improve inventory management. The supply planning team should revise the ABC(-XYZ) classification monthly and inventory policy settings every three months. As new products may be added to the core product portfolio each month after at least one year of sales data is available, it is important to update system performance regularly. Although replenishment settings are dynamic and respond to changing demand, recalculating them regularly is essential for two reasons: incorporating new data and establishing a company habit.

The simulation manual can be found in Appendix A.8. To ensure a sustainable solution, it is important to create an automated simulation template in SAP. Additionally, in future, the simulation could be partnered with SAP to directly use daily demand data.

An issue that we faced during this thesis was the time-consuming task of data cleansing. We recommend that PVM prioritize building an advanced data structure and invest in useful tools that can automate or significantly aid in the data cleansing process. This is necessary for achieving a reasonable level of data quality in an existing dataset while being practical and cost-effective.

6.2 Contributions

Through this thesis, we have made several contributions to PVM's inventory management. Specifically, we developed a stockout decision-making tool to assist the company in making informed operational decisions. We also provided a “mother-mapping” lookup document that connects data across departments, as well as a Power Query template for the ABC-XYZ classification, which has been integrated into PVM's hard disk. Lastly, we developed a simulation tool with an implementation manual to ensure the thesis’ deliverables are sustainable for the company.

Previous studies have explored the combination of multi-product, multi-location problems with product classification for fast-moving products. However, to our knowledge, no previous work has integrated customer differentiation into this framework. Additionally, considering transshipments and substitutions in the context of customer differentiation is a new aspect in inventory theory. Therefore, our contribution to the literature is a comprehensive study that aims to bridge this gap by employing simple analytics, simulation, and a straightforward greedy algorithm.

In our thesis, we distinguished two customer classes based on high and low priority. Low priority customers could find substitute products during stockouts, while high priority customers had order fulfillment priority but no substitute options. We explored a critical level policy to improve performance for high priority customers and reduce the total holding costs. However, the strict fill rate requirements along with the high sales volumes in the low priority class, greatly diminished the effectiveness of the critical level policy.

Moreover, we have observed that the partial pooling strategy performed well and effectively handled the excess inventory resulting from the minimum production quantity (MPQ).

6.3 Limitations and Future Research

Currently, the company's MPQ determination primarily relied on the production perspective, with limited input from stakeholders like planning, delivery and commerce. Future research could explore alternative approaches that take into account these stakeholders, aligning with business objectives and improving inventory management and cost optimization. While PVM is consistently striving to improve performance at its production facility, it is worth considering a thorough analysis of expansion opportunities. This analysis should evaluate the feasibility and potential benefits of increasing capacity at the existing factory or establishing a new facility in another region. With growing demand, these expansion options are crucial steps to be considered, and further research could provide valuable insights into their viability and implications.

For obvious reasons, this thesis has its limitations. These concern the supply constraints, demand distributions and data samples. The thesis aims to optimize the unconstrained performance of inventory management, but it is essential to acknowledge that supply chain issues, such as insufficient production capacity or material shortages (see Appendix A.9), prevent this state from being realized for a considerable period. Although these disruptions are significant in the real world, the system's flexibility to react to them through *ad hoc* decision-making is also greater than what could have been modeled (e.g., via proactive emergency shipments). Future research should explore this system's resilience in more depth. It is also worth noting that the results of this thesis are only applicable to products with normally distributed demands, commonly used for systems with high demand. For other types of known distributions, the selected inventory settings should be re-evaluated under similar experimental conditions. The data sample is limited to three years, and partially influenced by COVID-19 from June to December 2021. While we have already put significant efforts into data validation through sensitivity analysis, it would still be important to review the control parameters with new information in future. This thesis shows that the proposed greedy algorithm is effective, but future research should identify whether better alternatives can be used. Finally, it is interesting to consider the potential of safety stock in responding to global supply chain disruptions, particularly with the increasing frequency of natural disasters, catastrophes, and pandemics.

Bibliography

Agrawal, N., & Smith, S. A. (2015). Multi-location inventory models for retail supply chain management: A review of recent research. *International Series in Operations Research and Management Science*, *223*, 319–347.

Axsäter, S. (2015). *Inventory control* (2nd ed.). Lund, Sweden: Springer.

Axsäter, S., Howard, C., & Marklund, J. (2013). A distribution inventory model with transshipments from a support warehouse. *IIE Transactions (Institute of Industrial Engineers)*, *45*(3), 309–322.

Azimi, P., Ghanbari, M. R., & Mohammadi, H. (2012). Simulation modeling for analysis of a (Q, r) inventory system under supply disruption and customer differentiation with partial backordering. *Modelling and Simulation in Engineering*, *2012*.

Bartmann, D., & Bach, M. F. (2012). *Inventory control: models and methods* (Vol. 388). New York, NY: Springer Science & Business Media.

Chiadamrong, N., & Piyathanavong, V. (2017). Optimal design of supply chain network under uncertainty environment using hybrid analytical and simulation modeling approach. *Journal of Industrial Engineering International*, *13*(4), 465–478.

Chinello, E., Lee Herbert-Hansen, Z. N., & Khalid, W. (2020). Assessment of the impact of inventory optimization drivers in a multi-echelon supply chain: Case of a toy manufacturer. *Computers and Industrial Engineering*, *141*(106232), 1–22.

Chopra, S., & Meindl, P. (2016). *Supply chain management: Global edition*. *Supply Chain Management: Global Edition* (6th ed.). Harlow, England: Pearson Education Limited.

Christopher, M. (2000). The Agile Supply Chain. *Industrial Marketing Management*, *29*(1), 37–44.

Cohen, M. A., Kleindorfer, P. R., & Lee, H. L. (1988). Service constrained (s , S) inventory systems with priority demand classes and lost sales. *Management Science*, *34*(4), 482–499.

Cormen, T. H., Leiserson, C. E., Rivest, R. L., & Stein, C. (2007). *Introduction to algorithms*. *Bioinformatics: A Concept-Based Introduction* (3rd ed.). Cambridge, MA & London, England: The MIT Press.

Crittenden, A. B., Crittenden, V. L., & Crittenden, W. F. (2017). Industry Transformation via Channel Disruption. *Journal of Marketing Channels*, *24*(1–2), 13–26.

Drent, M., & Arts, J. (2020). Expediting in two-echelon spare parts inventory systems. *Manufacturing & Service Operations Management*.

Escalona, P., Angulo, A., Weston, J., Stegmaier, R., & Kauak, I. (2019). On the effect of two popular service-level measures on the design of a critical level policy for fast-moving items. *Computers and Operations Research*, *107*, 107–126.

Escalona, P., Ordóñez, F., & Kauak, I. (2017). Critical level rationing in inventory systems with continuously distributed demand. *OR Spectrum*, *39*(1), 273–301.

Garcia, F. A., Marchetta, M. G., Camargo, M., Morel, L., & Forradellas, R. Q. (2012). A Framework for Measuring Logistics Performance in the Wine Industry. *International Journal of Production Economics*, *135*(1), 284–298.

Göçken, M., & Dosdoğru, A. T. (2015). (R, s, S) inventory control policy and supplier selection in a two-echelon supply chain: An optimization via simulation approach. *Proceedings of the 2015 Winter Simulation Conference*, 2057–2067.

Gonçalves, J. N. C., Sameiro Carvalho, M., & Cortez, P. (2020). Operations research models and methods for safety stock determination: A review. *Operations Research Perspectives*, *7*(August), 100164.

Gümüs, A. T., & Güneri, A. F. (2007). Multi-echelon inventory management in supply chains with uncertain demand and lead times: Literature review from an operational research perspective. *Proceedings of the Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture*, *221*(10), 1553–1570.

Hai Nam, V. (2019). The Strategic Development in Logistics in Vietnam. *European Journal of Engineering Research and Science*, *4*(6), 69–73.

Hu, J., Watson, E., & Schneider, H. (2005). Approximate solutions for multi-location inventory systems with transshipments. *International Journal of Production Economics*, *97*(1), 31–43.

Jammernegg, W., & Reiner, G. (2007). Performance Improvement of Supply Chain Processes by Coordinated Inventory and Capacity Management. *International Journal of Production Economics*, *108*(1–2), 183–190.

Janssen, F., Heuts, R., & De Kok, T. (1998). On the (R, s, Q) inventory model when demand is modelled as a compound Bernoulli process. *European Journal of Operational Research*, *104*(3), 423–436.

Kranenburg, B., & van Houtum, G. J. (2009). A new partial pooling structure for spare parts networks. *European Journal of Operational Research*, *199*(3), 908–921.

Kranenburg, B., & Van Houtum, G. J. (2014). *Spare Parts Inventory Control under System Availability Constraints* (Vol. 227). Springer.

Kumari, A. G. K. C., Wijayanayake, A. N., & Niwunhella, D. H. H. (2022). Lateral transhipment inventory models : A systematic literature review of models and solution approaches. *18th International Conference on Business Management*, *18*, 54–80.

Law, A. (2014). *Simulation Modeling and Analysis* (5th ed.). New York, NY: McGraw-Hill Education.

Li, X., & Chen, Y. (2010). Impacts of supply disruptions and customer differentiation on a partial-backordering inventory system. *Simulation Modelling Practice and Theory*, *18*(5), 547–557.

Nahmias, S., & Demmy, W. S. (1981). Operating characteristics of an inventory system with rationing. *Management Science*, *27*(11), 1236–1245.

Najjartabar Bisheh, M., Nasiri, G. R., Esmaeili, E., Davoudpour, H., & Chang, S. I. (2022). A new supply chain distribution network design for two classes of customers using transfer recurrent neural network. *International Journal of System Assurance Engineering and Management*, *13*(5), 2604–2618.

Nakandala, D., Lau, H., & Shum, P. K. C. (2017). A lateral transshipment model for perishable inventory management. *International Journal of Production Research*, *55*(18), 5341–5354.

Núñez-Merino, M., Maqueira-Marín, J. M., Moyano-Fuentes, J., & Martínez-Jurado, P. J. (2020). Information and Digital Technologies of Industry 4.0 and Lean Supply Chain Management: A Systematic Literature Review. *International Journal of Production Research*, *39*(1), 5034–5061.

Paterson, C., Teunter, R., & Glazebrook, K. (2012). Enhanced lateral transshipments in a multi-location inventory system. *European Journal of Operational Research*, *221*(2), 317–327.

Paul, B., & Rajendran, C. (2011). Rationing mechanisms and inventory control-policy parameters for a divergent supply chain operating with lost sales and costs of review. *Computers and Operations Research*, *38*(8), 1117–1130.

Perfetti Van Melle. (2019). Our Mission & Vallues. Retrieved from https://www.perfettivanmelle.com/who-we-are/at-a-glance

Princes, E. (2020). Facing Disruptive Challenges in Supply Chain 4.0. *International Journal of Supply Chain Management*, *9*(4), 52–57.

Rădăşanu, A. C. (2016). Inventory Management, Service Level and Safety Stock. *Journal of Public Administration, Finance and Law*, *5*(9), 145–153.

Rai, S., Gross, E., & Ettam, R. K. (2015). Applied simulation and optimization. In M. M. Mota, I. F. D. La Mota, & D. G. Serrano (Eds.), *Applied simulation and optimization: In logistics, industrial and aeronautical practice* (Eds, pp. 277–301). Springer.

Robinson, S. (2004). *Simulation: The Practice of Model Development and Use*. Palgrave Macmillan.

Schmitt, A. J., Sun, S. A., Snyder, L. V., & Shen, Z. J. M. (2015). Centralization Versus Decentralization: Risk Pooling, Risk Diversification, and Supply Chain Disruptions. *Omega*, *52*, 201–212.

Scholz-Reiter, B., Heger, J., Meinecke, C., & Bergmann, J. (2012). Integration of demand forecasts in ABC-XYZ analysis: Practical investigation at an industrial company. *International Journal of Productivity and Performance Management*, *61*(4), 445–451.

Shu, J. (2010). An efficient greedy heuristic for warehouse-retailer network design optimization. *Transportation Science*, *44*(2), 183–192.

Silver, E. A., Pyke, D. F., & Thomas, D. H. T. (2017). *Inventory and production management in supply chains* (Fourth, Vol. 15). Boca Raton, FL: CRC Press.

Stojanović, M., & Regodić, D. (2017). The significance of the integrated multicriteria ABC-XYZ method for the inventory management process. *Acta Polytechnica Hungarica*, *14*(5), 29–48.

Tagaras, G. (1999). Pooling in multi-location periodic inventory distribution systems. *Omega*, *27*(1), 39–59.

Tempelmeier, H. (2006). Supply chain inventory optimization with two customer classes in discrete time. *European Journal of Operational Research*, *174*(1), 600–621.

Teunter, R. H., Babai, M. Z., & Syntetos, A. A. (2010). ABC Classification: Service levels and inventory costs. *Production and Operations Management*, *19*(3), 343–352.

Tlili, M., Moalla, M., & Campagne, J. P. (2012). The trans-shipment problem in a two-echelon, multi-location inventory system with lost sales. *International Journal of Production Research*, *50*(13), 3547–3559.

Van der Heide, G., Van Foreest, N. D., & Roodbergen, K. J. (2018). Optimizing stock levels for rental systems with a support warehouse and partial backordering. *European Journal of Operational Research*, *265*(1), 107–118.

Van der Heijden, M. (2020). Safety stocks: The (R, S) inventory control system [Powerpoints]. Retrieved from https://canvas.utwente.nl/courses/6952/files/1965536?module_item_id=189295

Veinott Jr., A. (1965). Optimal policy in a dynamic, single product, nonstationary inventory model with several demand classes. *Operations Research*, *13*(5), 761–778.

Vitasek, K. (2013). Supply Chain Management CSCMP Glossary. *Council of Supply Chain Management Professionals*.

Vo-Thai, H. C., Hong-Hue, T. H., & Tran, M. L. (2021). Technological- and Non-Technological Innovation During the Growth Phase of Industry Life Cycle: An Evidence From Vietnamese Manufacturing Enterprises. *SAGE Open*, *11*(3).

Vo, T. T., & Nguyen, A. D. (2012). Experiences of Vietnam in FDI Promotion: Some Lessons for Myanmar. *Economic Reforms in Myanmar: Pathways and Prospects*, (10), 131–172.

Wunker, R. (1994). The Laws of Vietnam Affecting Foreign Investment. *International Lawyer*, *28*(2), 363.

Appendices

Appendix A.1 Infographic Timeline

Fig. A.1 shows an infographic overview of PVM Vietnam over the past 25 year, of which essential milestones are:

- the company entering Vietnam as a joint venture in April 1997;
- the start as an independent enterprise in May 2002;
- operationalizing the Bình Dương factory in September 2006;
- the opening of the local Ho Chi Minh City headquarters in 2015.

Fig. A.1 Infographic timeline of Perfetti Van Melle Vietnam.

Appendix A.2 Snapshot Stock Availability Dashboard

Fig. A.2 depicts a snapshot of the stock availability dashboard created in Power BI. On the left, it offers the opportunity to filter on products and brands. There it is also possible to change the period and required inventory levels. A visualization of the stock levels, demand and stockouts is provided in the center of the sheet, supported by performance data in the upper right corner.

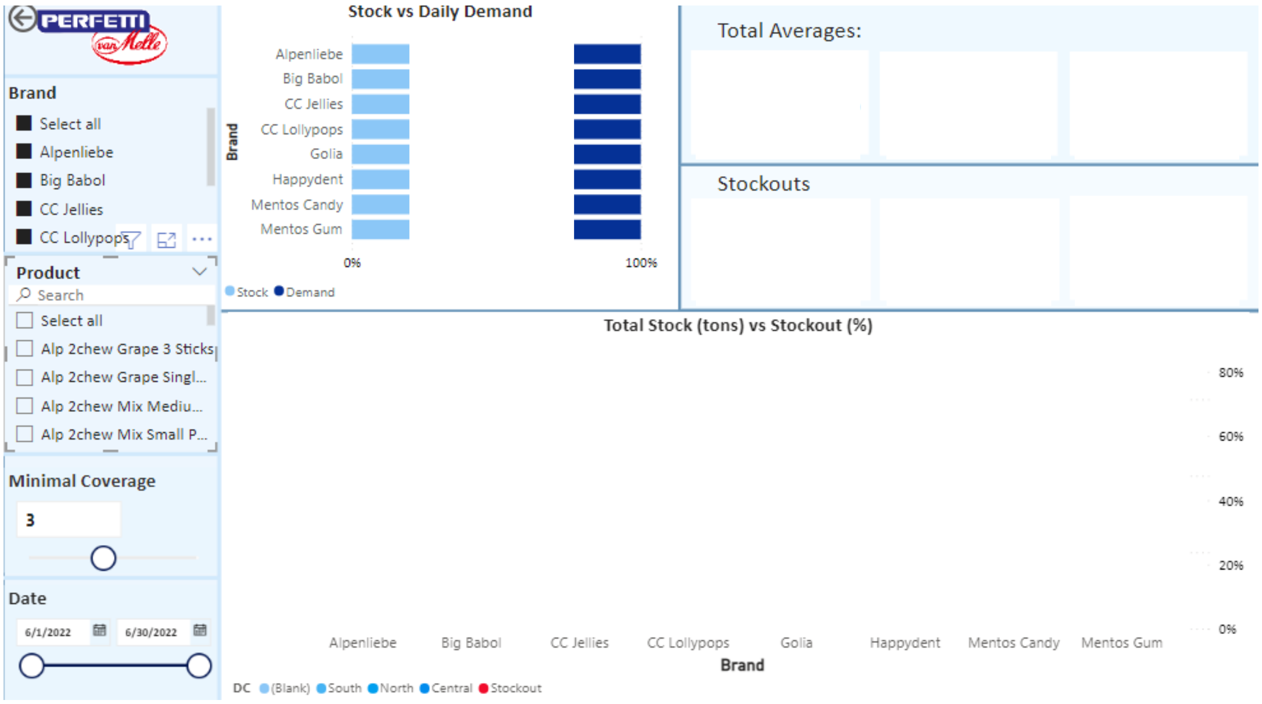

Fig. A.2 Stock availability report in Power BI.

Appendix A.3 Stock Availability Report Manual

This appendix presents the manual to make the stock availability report, which serves as input for the S&OP meetings.

Appendix A.4 Link Source Data by Mother Mapping.

Fig. A.3 shows the procedure of connecting different source files (input) based on their product code and corresponding mother map (link) to get the desired outcome (output).

Sales	
Code	**Volume (tons)**
827241	10.87
876201	0.05
827241	0.91
872943	2.61

Finance	
Code	**COGS/ton**
872943	134.086 VND
872653	47.697 VND

Input

Lookup Table	
Mother map	**Code**
BB Sploosh Cup	827241
BB Sploosh Cup	872943
CC Sour Belt 1m	872653
CC Sour Belt 1m	876201

Link

Output Table		
Mother map	**Volume (tons)**	**COGS/ton**
BB Sploosh Cup	14.39	134.086 VND
CC Sour Belt 1m	0.05	47.697 VND

Output

Fig. A.3 Fictious presentation of the mother mapping process.

Appendix A.5 Simulated Demand

Figs. A.4 and A.5 plot the detrended demand generated in the simulation via the normal distribution and historical data. In general, both demands follow about the same pattern.

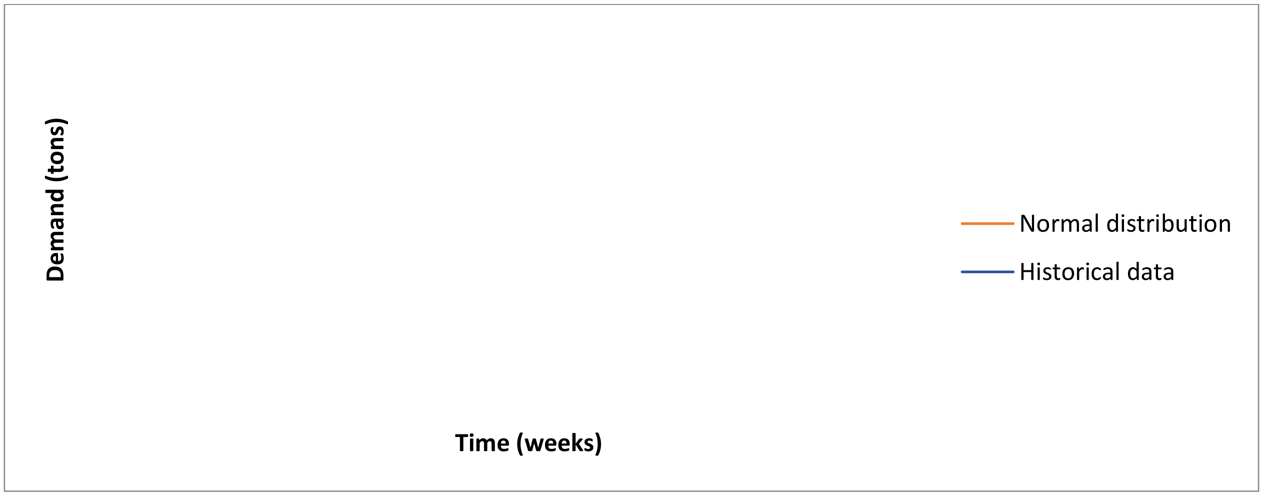

Fig. A.4 *Comparison of the simulated demand: normal distribution vs. historical data.*

The normally distributed data has been compared to the historical data using a chi-square test. The normal distribution (310, 2.59) passed the chi-square test with significance of 5% and 17 degrees of freedom.

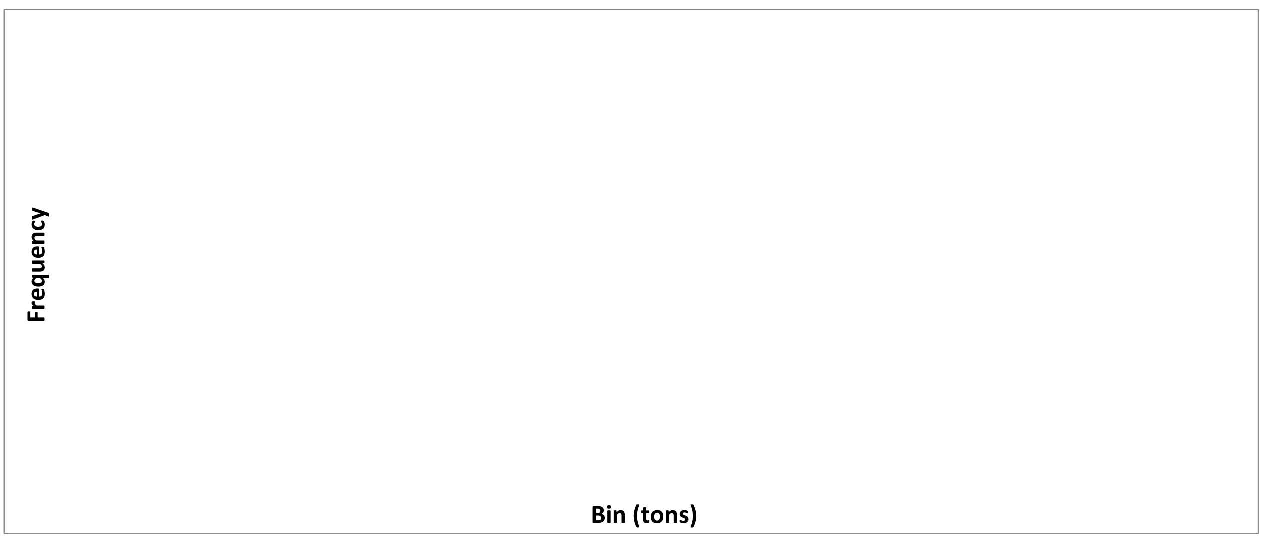

Fig. A.5 *Chi-square test on historical and normally distributed demand.*

Appendix A.6 Effect of the Minimum Production Quantity

This appendix provides the order-up-to-levels for the SDC including the MPQs during the optimization. The algorithm of the second stage proposed values for S provided in Table A.1, and the final optimization led to Table A.2. From Table A.1 it is visible that the values follow the same trend that as the analytical approximated ones, but that the levels are generally lower (see Section 5.1.3). Especially for the southern DC and X-class, the levels are considerably lower.

The fill rates presented in Table A.3 indicate that the overall fill rate was identical to those of the ABC-XYZ listed given in Section 5.1.3. Additionally, the differences in costs were negligent, as the holding cost just slightly increased.

Table A.1 Order-up-to-level per class and Distribution Center by the algorithmic approach (stage 2).

Class	AX	AY	AZ	BX	BY	BZ	CX	CY	CZ
SDC	XX	XX	XX	XX	XX	XX	XX	XX	XX
NDC	XX	XX	XX	XX	XX	XX	XX	XX	XX
CDC	XX	XX	XX	XX	XX	XX	XX	XX	XX

Introducing the critical level policy did not lead to significant changes, as the costs and MT fill rate improved hardly. The critical levels were only applicable for the CX class and S was never reduced, as either not all fill rate constraints were satisfied or costs increased.

Table A.2 Critical levels per class and Distribution Center by the algorithmic approach (stage 3).

Class	AX	AY	AZ	BX	BY	BZ	CX	CY	CZ
SDC	XX	XX	XX	XX	XX	XX	XX	XX	XX
NDC	XX	XX	XX	XX	XX	XX	XX	XX	XX
CDC	XX	XX	XX	XX	XX	XX	XX	XX	XX

Table A.3 DC fill rates under optimization stages 2 and 3.

	Total	SDC	NDC	CDC
Optimization stage 2	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X% (XX.X%)
Optimization stage 3	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X% (XX.X%)	XX.X% (XX.X%)

Table A.4 Channel fill rates under optimization stages 2 and 3.

	Traditional Trade	Modern Trade
Optimization stage 2	XX.X% (XX.X%)	XX.X%
Optimization stage 3	XX.X% (XX.X%)	XX.X%

Table A.5 Costs (in millions) under optimization stages 2 and 3.

	Inventory costs	Holding cost	Holding cost
Optimization stage 2	VND XX,XXX	VND XX,XXX	VND XX,XXX
Optimization stage 3	VND XX,XXX	VND XX,XXX	VND XX,XXX

Appendix A.7 Paired-T Approach

We applied a paired-t approach to analyze whether there was an improved performance in stage 3 compared to stage 2. For this tests we used a 95% confidence interval, from which we could not tell whether there was indeed an improved performance as the bounds were [-XX,XXX, XX,XXX] M, meaning zero was in it.

Table A.6 Confidence interval bounds (in millions) of the t-paired test.

Lower bound	Upper bound
VND -XX,XXX	VND XX,XXX

Table A.7 Inventory costs (in millions) per experiment trial of the paired-t test.

Trial	Stage 2	Stage 3
1	VND XX,XXX	VND XX,XXX
2	VND XX,XXX	VND XX,XXX
3	VND XX,XXX	VND XX,XXX
4	VND XX,XXX	VND XX,XXX
5	VND XX,XXX	VND XX,XXX

Appendix A.8 Simulation Manual

This appendix contains the manual provided to PVM to enable them to integrate the solution into their routine.

Appendix A.9 Tactical Level Ishikawa Diagram

Fig. A.6 shows the fishbone diagram to identify the current causes of inefficiencies in inventory management.

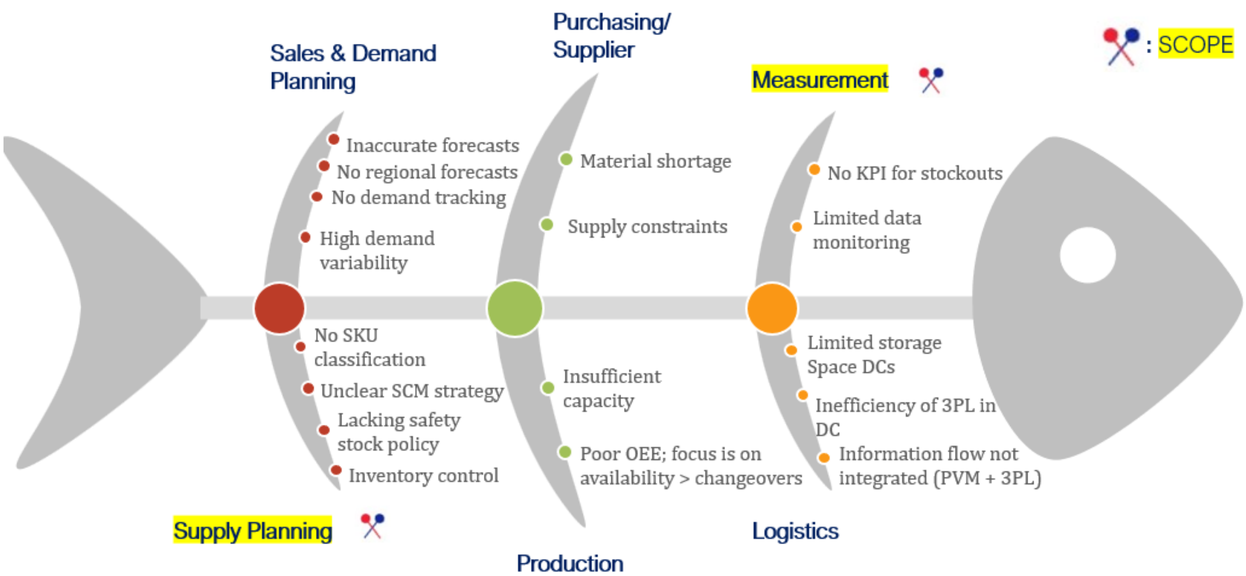

Fig. A.6 Ishikawa diagram representing the causes of inefficiencies in inventory management