Choosing between regional and central warehouses for different items: a case study in grocery retail

Miguel Ruano Neto Veiga de Macedo

Master's Dissertation

University Supervisor: Gonçalo Figueira

U.PORTO

FEUP FACULDADE DE ENGENHARIA UNIVERSIDADE DO PORTO

Mestrado Integrado em Engenharia e Gestão Industrial

2021-06-28

Abstract

Often, retailers have to make decisions on where they hold inventory in their internal supply-chain. This is a fundamental problem, as these decisions need to be revised periodically and affect costs across the entire operation The current work was developed in the context of a major grocery retailer and focused on the regional operation of the company in the Madeira Archipelago. The objective consists of developing an holistic methodology to decide, based on sales forecasts for each product inside the scope, whether to hold inventory on the regional warehouse in Madeira (defined as centralized flow) or on the mainland warehouse (defined as non-centralized flow).

The approach used consisted of (i) a cost modelling phase that predicted, for each product, the supply chain costs of using either flow and (ii) a prioritization phase, where the actual allocation to the regional warehouse was decided, based on savings per product and the of capacity constraint of the regional warehouse. The first phase involved an approximation of the retailer's replenishment policy by a standard (R,S) policy and the cost modelling of all relevant logistical operations. In the second phase, where the problem took the form of a classical optimization problem (bi-dimensional Knapsack Problem), a simple greedy algorithm was used to generate a reasonably good allocation proposal.

The main goal of the project was achieved, as a system was developed that successfully captured all relevant costs in the supply chain, suggesting an allocation that reduces total supply chain cost. Although these should be subject to further study due to deviations in the model, promising savings of about 40000 € or 6% of the analysed supply chain costs of resulting from supplying these products during the two month study period were identified. Important insights on the nature of the operation were also achieved. The relative large importance of the transportation costs in the operation, due to the maritime transport to Madeira, and the sharp diminishing returns from the sequential centralization of products might have a role in simplifying the allocation task.

This produced work will become the basis for the creation of an allocation tool that the retailer will use in the future to periodically manage the allocation process.

Resumo

Frequentemente, retalhistas têm que tomar decisões sobre onde armazenar inventário na sua cadeia de abastecimento interna. Este é um problema de grande importância, visto que estas decisões carecem de revisão periódica e afetam diversos custos de operação em toda a cadeia de abastecimento. O trabalho aqui apresentado foi realizado no contexto de um grande retalhista nacional, com foco na operação desse retalhista na Região Autónoma da Madeira. O objetivo do projeto foi o desenvolvimento de uma metodologia holística que apoiasse a seleção dos produtos a armazenar no armazém regional na Madeira (denominados de centralizados) e dos produtos a armazenar no armazém principal, em Portugal continental (denominados não-centralizados). A abordagem tomada consistiu numa fase (i) de modelação de custos da cadeia de abastecimento em cada um dos fluxos, para cada produto analisado e (ii) uma fase de alocação onde a decisão sobre o fluxo de cada produto teve por base as poupanças que a respectiva alocação ao armazém regional traria, bem como a capacidade limitada desse mesmo armazém. A primeira fase teve por base uma aproximação da política de abastecimento do retalhista a uma política (R,S) standard e a modelação de custos de todas as operações logísticos relevantes. Na segunda fase, onde o problema tomou a forma de um problema de otimização clássico (problema de *Knapsack* nbi-dimensional), um algoritmo ganancioso simples foi empregue para gerar uma proposta de alocação aproximada à alocação ótima. O principal objetivo do projeto foi atingido, dado que o sistema desenvolvido captura com sucesso todos os custos de cadeia relevantes, sugerindo uma alocação que reduz o custo de cadeia total. Apesar de ainda carecerem de estudo adicional, devido aos desvios de previsão apresentados pelo modelo, poupanças promissoras de cerca de 40000€ ou 6% nos custos de cadeia associados aos produtos em estudo poderiam ter sido capturadas durante o período alvo, se a alocação sugerida pelo modelo tivesse sido utilizada. Outro resultado importante do projeto foram aprendizagens cruciais sobre a natureza da operação logística regional. A primeira conclusão, foi o grande peso relativo dos custos de transporte na operação da Madeira, devido sobretudo ao custo do transporte marítimo. A outra aprendizagem importante foi o reconhecimento do rápido declínio em ganhos resultantes da alocação sucessiva de itens ao armazém regional, que poderá simplificar a tarefa de alocação no futuro. O trabalho realizado será a base para o desenvolvimento de uma ferramenta de gestão do processo de alocação de produtos ao armazém regional, a ser desenvolvida pelo retalhista.

Acknowledgments

I would like to thank my supervisors João Almeida and Gonçalo Figueira, for all of the valuable discussions and guidance provided during the project. Your advises, time and comments are the unseen pillars of this thesis!

I also want to thank SONAE MC for having received me during this project, and especially Rui Braz for putting his trust in a random student that was in need of an internship! This thank you extends to the people in SONAE that took their time to help me with this project: Paulo Pereira, Paulo Pinho, Tânia Duarte and Luís Homet! Even though I never got to meet some of you in person due to this forsaken pandemic, your "remote" help is present all throughout this work (and i owe each one a beer!)

This thesis is also dedicated to some of the professors that shaped me most in these FEUP years, especially Ana Camanho, Henriqueta Nóvoa and José Faria.

To my parents, my sister Rita and my brother João who are way more cooler than me, even though I don't usually admit it. To my grandma Maria Antónia, who I love very very much and who gives me shelter and food every time I need it. (Obrigado Avó Tonita!)

And finally to my friends from close and far away, how are always ready to go on adventures! (With a special mention to Teresa, who is always in my heart!)

I really had a blast these 6 years! On to new chapters!

"What is the real goal? Nobody here has even asked anything that basic."

Eliyahu M. Goldratt

Contents

Acronyms and Symbols

DC	Distribution Center
SKU	Stock Keeping Unit
PBL	Picking-by-line
PBS	Picking-by-store
MVP	Minimum viable product
BRE	Back Room Effect
SS	Safety Stock
CS	Cycle Stock
MDQ	Minimum Display Quantity
ESPRC	Expected shortages per replenishment cycle

List of Figures

List of Tables

Chapter 1

Introduction

This introductory chapter starts with explaining the motivation to undertake the project (1.1), follows with a short description of the approach used to tackle the problem (1.2), and closes with an overview of the complete document, in order to give an overall picture of the accomplished work (1.3).

1.1 Motivation

Modern retail companies co-exist with the constant need to balance costumer satisfaction and product availability at stores with the internal push to run an operation that is as efficient as possible in order to increase operating margins. In fact an adaptable and efficiently managed supply chain as become a competitive advantage in the grocery retail market. These supply-chains often involve multiple stakeholders spread out geographically, which are increasingly hard to manage efficiently as a whole. This happens because of the shear amount of operational, tactical and strategic decisions that need to be regularly taken and implemented just to keep the supply-chain aligned with the ever changing costumer demands. One of such decisions that retailers need to make is where to hold inventory of a specific product in their internal supply chain. Retail companies have a network of warehouses with different typologies in different locations, providing specific advantages and disadvantages for the storage of products with certain characteristics and demand patterns. Concurrently, retailers need to review these decisions in a timely manner, to respond to the ever-increasing dynamic of the consumer demand.

These decisions become even more challenging when they involve unfavorable geography, like supplying stores in islands because these operations are subject to several specific constraints. A limited selection of local suppliers, impossibility of sourcing some products locally, very stretched lead-times, heavy costs related with transportation (when compared with mainland locations) are some of these special limitations. Retailers struggle to offer a similar service to these challenging locations as the one experienced in the mainland, whilst maintaining profitability. This calls for close attention to decisions that affect the total efficiency of the operation.

Ultimately, retailers should build processes for reviewing product allocation to warehouses, which take into account an holistic approach to all the challenges and costs derived from operating in regions like this. They also ought to have a regular enough periodic review so as to not be outpaced by changing consumer trends.

1.2 Objectives and Methodology

The company's project was initiated by the Logistics Innovation Department, and concerns one such problem as specified above. It has the goal of revamping the previous process of SKU flow selection in the autonomous region of Madeira, in Portugal. More specifically, the decision concerns whether to hold stock of a certain *Stock Keeping Unit* (SKU), sold in Madeira, on the island's regional warehouse or to ship the stock already sorted to the stores from the mainland warehouses.

The project involved the study of new decision criteria that are at the core of the selection process, by including other input variables and proposing a new selection algorithm, more fitting to the company's objectives. These were the reduction of total supply chain cost in the region while also attending to commercial goals and operational restrictions. The final objective was to use this knowledge to later build a Decision Support System and to acquire a deep understanding of the interactions and consequences of both possible product allocations.

The project's key stakeholder was the supply-chain management department, which is the direct responsible for the allocation process. Nonetheless, every team involved with the retailer's regional supply chain can also be considered a stakeholder. After an initial period of getting familiar with the process, defining the scope and exact goals of the project, two large phases emerged as core to the approach. The first was evaluating all relevant costs that the product allocation could entail into two holistic cost functions (one for each possible allocation) and the subsequent calculation of the result of these functions for each SKU. The second phase was the sequential assignment of the products based on the outputs of the first phase and the relevant operational constraints. This generated a recommendation for the allocation. This two-phase allocation approach was encapsulated into a prototype version of what would later originate the final tool to be used by the key stakeholder.

1.3 Dissertation Structure

The structure of the dissertation is as follows. In Chapter 2, the problem is defined in depth, listing how all relevant stakeholders and supply chain stages are affected by the SKU allocation. In Chapter 3, the existing knowledge on product allocation problems present in the current literature is summarized, as well as other relevant topics for the modelling phase such as the basics of inventory replenishment. In Chapter 4, the two phase approach is thoroughly explained, with the used models being defined and the individual cost proxies being focused. Chapter 6 presents the results of the develop method for the studied retail company, as well as an analysis of model deviations. Chapter 7, comments specifically on the produced results and includes a broader reflection on the work done. This chapter closes with main conclusions and the future enhancements to the tool and approach resulting from the project.

Chapter 2

Problem Definition

Chapter 2 offers a complete view of the SKU allocation problem. It does this by first exploring how these decisions affect each stakeholder in the company's internal supply chain (2.1) and then by providing a first look to the steps taken during the project's implementation (2.2).

2.1 Problem context

Retailers tend to operate complex supply chains comprised of a network of warehouses and distribution centers (DCs) of different typologies, spread across their geographical area of action, in order to achieve the intended levels of service to the stores while running operations as cost efficiently as possible.

These locations preform several tasks like order consolidation, picking operations and inventory storage. This last point poses challenges to retailers, when it comes to deciding at which locations to hold inventory of each specific SKU. This decision impacts operations across the complete supply chain, in very different ways, which makes this problem non-trivial.

Complexity increases drastically when one considers that these decisions need to be done for tens of thousands of different products, and that any proposed allocation solution must account for constrained warehouses, which, as Holzapfel et al. (2018) show, might result in solutions that greatly differ from those of generally unconstrained problems.

SKU allocation decisions need also to be frequently revised, because they are partly influenced by product demand. Naturally, shifting consumer preferences, introduction of new products or the natural seasonality of some SKUs dictate that an efficient allocation solution today might not be adequate a few months ahead.

This point demands some agility from retailers. Flow alteration tasks take some time to implement, forcing retailers to preform changes weeks before demand patterns actually change, so as not to be caught off-guard with a non adjusted operation.

Holzapfel et al. (2018) point that supply-chain managers often tend to resort to simplistic rules of thumb in order to periodically perform the SKU allocation, which might lead to local optima

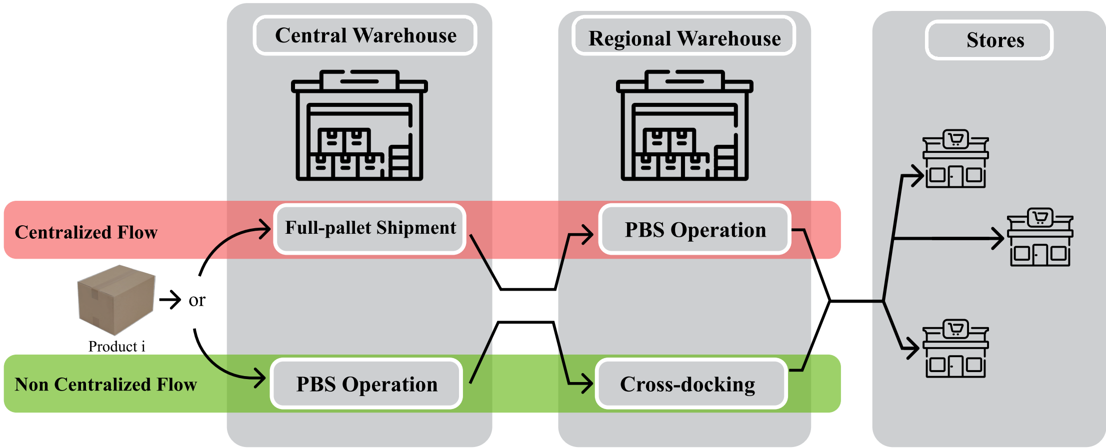

Figure 2.1: Nodes of the modeled supply chain

in certain nodes of the supply chain, but not necessarily originate the overall most efficient supply chain operation.

This master thesis aims at tackling all the points noted above in a coherent and holistic way. The goal is to reliably model the cost impacts of assigning each of the relevant SKU's to each of two possible locations in the supply-chain of a Portuguese retailer, based on intrinsic product characteristics and mid-term sales forecasts. The model produced is later used as the basis for a decision support system, which will periodically recommend the ideal product flow for each SKU included in the scope of the analysis. Along side with this thesis a Minimum Viable Product (MVP) for this tool was developed, which is the source of results studied in Chapter 6.

Supply chain structure

Even though some standard costs impacts of SKU allocation can be generalized, as summarized in the existing literature, the specific costs drivers are always a consequence of the structure of the supply chain in question.

This case-study focuses on a central warehouse located in mainland Portugal, a regional warehouse in the Madeira autonomous region and the retailers stores in the archipelago, as seen in Figure 2.1.

Both the central and the regional warehouses hold stocks of products and run Picking-by-line (PBL) and Picking-by-store (PBS) order picking operations to replenish the stores of the region, although capacity of the Madeira warehouse is fairly limited.

This thesis concerns itself with the products sold in Madeira that are replenished using the PBS flow and that are sourced from suppliers in the mainland. The company handles logistics of these products internally, consolidating all product flows to the archipelago in its mainland warehouse, in order to achieve transport economies of scale in the very costly maritime shipment. Containers are periodically shipped to the island departing from the central warehouse.

For each SKU of this group, the company assigns one of two possible flows:

- **Centralized -** In this flow, store orders originate from the regional warehouse. Some quantity of inventory is held permanently at the regional warehouse that acts as the supplier for the stores in the region. The orders for the stores are prepared through a PBS operation. Bulk replenishment to the regional warehouse also originates from the stock held at the central warehouse, only in single product pallet formats (i.e. in the same format that they are received form the supplier).
- **Non-Centralized -** In this flow, store orders originate from the mainland warehouse. No inventory of the product is held at the regional warehouse. Instead, store orders for the region are picked in the PBS operation in the mainland warehouse and are already shipped in store-ready pallets. In this case the regional warehouse is solely used as a transshipment platform.

The retailer opts to assign each product to only one flow, to avoid task duplication and reduce complexity in the replenishment process.

An holistic view is needed in order to define the optimal quantity, which requires an understanding of how SKU allocation to one flow or another might influence the entire chain, starting with the warehouses, followed by transportation and stores as seen in the scheme.

Warehousing

Product allocation naturally affects warehousing costs. Relevant operations that need to be taken into account are picking operations (namely PBS) and cross-docking operations. It is also important to mentions the role of inventory opportunity costs.

A PBS is a product flow type that at its core deals with holding some stock, somewhere in the internal supply-chain, in the same warehouse where store-orders are picked. Items are stored at defined positions and the picking is done by collecting products from their position, when a store places an order. Then, the worker travels around the storage racks, progressively picking all the items that make up the store's order, in the specified quantities. This is a form of *pick-to-order* (Rushton et al. (2014)).

When considering cost impacts of PBS, it is also crucial to account for provisioning tasks, like receiving bulk stock and storing it in the racks, which necessarily precede the picking itself. According to Rushton et al. (2014), because there's an incentive to keep picking areas as small as possible to increase picking efficiency, it is typical to separate the reserve inventory and the inventory that is actually being used in the picking operation. As such, the warehousing system needs to assign picking and reserve locations to individual items. This separation between reserve and picking inventory can be vertical (i.e. using multi-level pallet storage racks), with picking operations using the ground-floor rack positions and reserve stock being placed at the higher racking levels. Note that in smaller warehouses, the available number of racks might represent a constraint to the range of products held at the location, with limitations originating either in the usage of picking locations or in the usage of reserve locations. Figure 2.2 is a visual representation of all the points made in this paragraph.

Figure 2.2: An example of a PBS operation in Retail

In cross-docking, goods are received at a warehouse and dispatched without storing them as inventory first. Goods can be effectively just moved from the incoming bay to the dispatching bay (Rushton et al. (2014)). In this supply-chain, PBS pallets already assigned to stores arrive bundled form the mainland in the same shipping containers and are then distributed to the trucks that take them to their final destination at the region's stores.

Transportation

Efficiency of transportation is a major driver of supply chain costs for retailers. There is a constant effort to reduce costs by maximizing volumes transported in each truck or shipping container, and thus reducing the total number of truckloads needed to support the operation. The number of shipped boxes can be increased by utilizing more volume in each pallet. This in turn can be achieved by packing boxes to a taller height or by a more efficient box stacking, which reduces the volume wasted in empty space between boxes. For the supply-chain in study, this factor is even more relevant given the very high costs of shipping containers from mainland Portugal to Madeira. SKU allocation will ultimately affect transportation costs, because products stocked in the regional warehouse are shipped in single product pallets originating from the supplier, which tend to have the best space utilization possible. All other products are shipped bundled with boxes of other PBS items, necessarily resulting in lower volume utilization.

Stores

Costs incurred in stores can be related to shelf replenishment, inventory costs and stock-outs. All of these are a direct consequence of the amount of inventory held at the stores and the size of the orders received. These factors are in turn directly dependent on the orders' lead times and frequencies dictated by the SKU's warehouse allocation.

When incoming pallets arrive at the store, they are off-loaded from the truck and products are place in the respective shelves until these are full. It is possible that the number of units delivered in an order largely surpasses the available shelf space, and thus some of it needs to be held at the backroom of the stores. (Sternbeck and Kuhn (2014)) The results are adverse consequences, namely additional tasks and movements associated with the already costly task of shelf replenishment, due to the need to preform this task several times between order arrivals.

The level of stock held in stores is the source of inventory costs. This changes for each flow, as the average stock on-hand at the stores will also be different. The money tied-up in stock on-hand is reflected in the retailer's inventory opportunity cost, by multiplying the amount of inventory held by the given rate of capital opportunity cost for the target period.

Finally, stock-out costs try to capture the loss of sales that did not happen because stock available at the store shelve could not satisfy demand. Stock-outs can have far more overreaching impacts, specially in loss of costumer goodwill, that retailers need to keep in mind when making allocation decisions. These were however left out of the cost model. Stock-outs at stores can be divided into two main categories: (i) stock-outs with no inventory, where there is not any inventory of the product at the store; (ii) and stock-outs with inventory, where there is extra inventory in the backroom of the store that for some reason is not put in display on the shelf. The root causes of both stock-out types are very hard to pin-point as they can have many different causes, many of them influenced by SKU allocation.

Stock-outs with no inventory might relate to SKU allocation as this affects order lead times and thus the speed with which the company responds to spikes in demand. Additionally, there might be differences of service level (i.e the proportion of times orders placed by the store to the warehouse are accurately fulfilled) between the regional and the main warehouse.

Stock-outs with inventory can be influenced by SKU allocation because of the backroom effect, described by Eroglu et al. (2013). Non-centralized SKUs require higher average stocks to be held at stores due to larger lead times and order windows. As the total average inventory in the backroom of the stores increases, the probability of stock-outs with inventory might also increase, as stores backrooms become more full and shelf replenishment between orders becomes harder to perform . A case could be made for minimizing backroom inventories by an effective selection of SKUs that are centralized, thus minimizing this type of stock-outs not only for the centralized SKUs, but for the whole range sold at the store. It should be said that this relation is not as solid as other cost impacts.

Shrinkage

Shrinkage costs are associated with loss of value of stocked products. They are the result of incorrect handling that damages the product or by the product surpassing its expiration date. Price markdowns on products nearing the expiration date, in order to encourage their purchase by consumers - a practice quite common in fresh products - are also counted as shrinkage costs. For sake

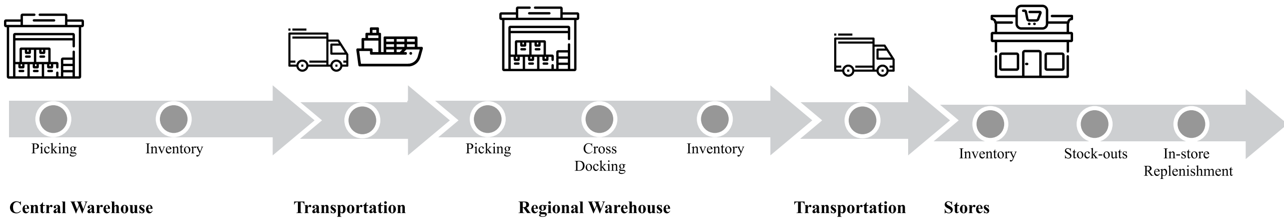

Figure 2.3: Stakeholders impacted by the SKU allocation

of simplicity, costs related with shrinkage were left out of the scope of the problem. The reason was twofold: (i) fresh product categories, which are much more sensitive to shrinkage were not included in the analysis and (ii) potential increases in shrinkage due to additional handling operations at the regional warehouse (in the centralized flow) were not considered to greatly increase the level of shrinkage due to mishandling.

2.2 Methodology

Figure 2.3 sums-up all the mentioned cost drivers, acting as the base for the work in the project described in the following chapters. Starting from this foundational mapping of the supply chains costs, several defined steps were done to bring it into a functioning practical application. This implementation was comprised of 3 large phases, that are outlined in Table 2.1

In the first phase the focus was on understanding the current state of the process with the stakeholders and defining the problem scope precisely and in an actionable way. Both the Centralized and Non-Centralized flows were thoroughly mapped, and costs that were deemed as relevant were aggregated together to form a first draft of the cost modelling function.

In the second phase, the problem was broken down into several cost components and precise cost drivers were progressively defined working with the stakeholders. At this point, some costs initially considered relevant were dropped, as a result of further study of historical data. The creation of the proxies meant a strong focus on the data treatment, which translated to the creation of extensive data extraction and cleaning pipelines for each cost proxy. This required special treatment of outliers and missing values to keep the model coherent and rigorous. This ended as all major cost proxies were created and joined together in the cost modelling tool, generating a prediction of total supply chain costs for each product in each flow.

In the final third phase, the product allocation heuristic was developed, providing the final deliverable that the retailer expects from the model. Additional work was done to translate the suggested allocations into an actionable recommendation for the stakeholders that will use the tool. Estimated gains with allocation change, as well as a study of model deviations, were also an output of this phase.

Table 2.1: Methodology Review

Phase 1 Process mapping	Process mapping of both centralized and non-centralized product flows.
	Investigation of previous process for SKU allocation.
	Identification of groups of costs to include in the cost function; First draft of the cost function.
Phase 2 Cost-modelling	Data extraction, cleaning and transformation.
	Iterative creation of detailed cost drivers inside of each cost group.
	Further study on impact/relevance of certain cost factors.
	Accurate total supply chain cost for each product in both flows.
Phase 3 Item Allocation	Prioritization of allocation based on usage of constrained storage resources.
	Translation of generated solution into actionable allocation alteration procedures.
	Study of economical gains of allocation and of the model's prediction errors.

Chapter 3

Literature Review

With this chapter the objective is to induce the reader to the most important topics that lay the groundwork for the project. The literature review is organized as follows: Section 3.1 takes a broad look at the role of well-run supply chains as a competitive advantage in retail. Section 3.2 sheds light on the internal functioning of warehousing operations, while section 3.3 describes at length the importance and nature of SKU allocation tasks and presents the most up-to-date approaches to the problem. Section 3.4 focuses on the effect of order quantities on in-store logistics, also denominated as *Backroom Effect*. Section 3.5 explores the inner-works of inventory replenishment policies, that are used as one of the basis for the modelling done during the thesis. Finally, section 3.6 briefly presents the Knapsack Problem, as this will be relevant the final phase of item allocation.

3.1 Retail Supply Chains

Different definitions of the concepts of supply chain and logistics are abundant in the available literature, given the broad set of activities these terms encompass. Lapinskaitė (2014) sums up existing definitions, describing supply chain as *"the overall processes of the product movement starting from the supplier to the end customer"*. Rushton et al. (2014) refers to Logistics as *"the efficient transfer of goods from the source of supply through the place of manufacture to the point of consumption in a cost-effective way while providing an acceptable service to the customer"*.

The importance and difficulty of these operations often goes unnoticed by final costumers. They are characterised by complex interconnections that dictate that optimizing such systems is challenging. Rushton et al. (2014) note that it is the integration of the various logistical components that enables system wide optimums, focusing in the concept of *Total Logistics*, prioritizing the need to identify the opportunities and needed trade-offs of the complete system. Fernie and Sparks (2018) note that managing logistical operations wrongly can be extremely costly, due to warehousing, stock, transport and obsolescence costs. At the same time, well run logistical operations can provide fresher, higher quality products to final costumers while simultaneously reducing stock-outs and overall cost, thus becoming a competitive advantage for retailers. For this to happen, retailers must concern themselves with both information and product flows, reacting rapidly to alterations in costumer demand.

3.2 Warehousing Activities

A general understanding of warehousing operations is required to address the SKU assignment problem.

According to de Koster et al. (2007) warehouses take an important role in achieving the retail company's missions, like dealing with demand uncertainty, enabling transport economies of scale and order preparation. They are a tool to *"Accomplishing least total cost logistics commensurate with a desired level of customer service."*

Tompkins et al. (2010) identifies the sequential tasks that may take place in a warehousing operation, depending on a product's specific flow. These include receiving the products from the suppliers, which can either be followed directly by reshipment to a further destination (equating to a cross-docking operation) or by reserve storing, followed by a form of order picking, then sorting, packing and finally shipping. A scheme of possible warehousing flows is displayed in Figure 3.1.

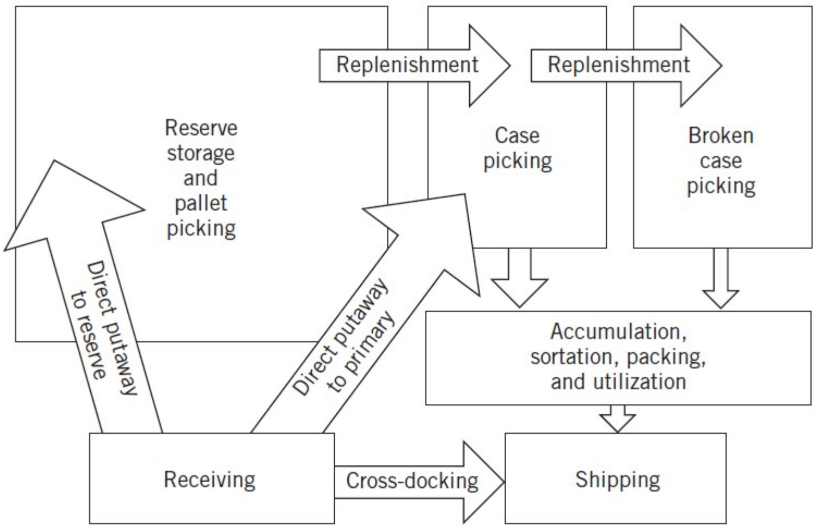

Figure 3.1: Warehousing Activities Scheme, adapted from Tompkins et al. (2010)

One of the key tasks of warehouses is order preparation. In retail, this means the consolidation of a variety of products required by a specific store into a single shipment. These operations are still largely manual, being a big source of warehouse labour costs. (Rushton et al. (2014))

Warehouses can choose to employ many different strategies for order preparation. One of these strategies is defined by Rushton et al. (2014) as *Pick-to-order* , where a worker travels through the warehouse, visiting the stock locations of each product and picking items until the order is complete. The *Pick-by-Store* method described in the previous chapter works in this fashion.

3.3 SKU-Warehouse Allocation Problem

An important decision that retailers face when dealing with their complex distribution networks, is deciding which articles to hold stock of in each warehouse in order to run the most optimized operation possible while attending to network constraints and the constant changing nature of consumer needs. This is the generic SKU allocation problem formulation

In defining a framework for supply-chain planning in grocery retailing, Hübner et al. (2013) identify a group of tasks related to product segmentation and allocation as one of the mid-term concerns of grocery retail supply-chain managers. In this group of tasks, product-warehouse-outlets assignment is included. It is referred that these decisions remain static for some time to alleviate capacity decisions. The authors note that stock can be held at central and regional warehouses but also that stores backrooms and shelf space are inventory holding locations. Managers should consider both the market pressures and the complex trade-offs between costs of inventories, shipment frequencies and handling activities when planning the appropriate level of centralization of a product.

Lovell et al. (2005) investigates the impact of product characteristics as a criteria for segmentation between different supply chain flows, in a European electronic products retailer. The authors focus on whether an item should be stored at the central European warehouses (centralized) or if it should be held at smaller, country level warehouses closer to the end consumer (de-centralized). Focus is given to product-value density, product turnover and defined service level. High value density favors more centralization, due to better inventory pooling effects and more efficient transportation. Higher turn-over products are said to favor a decentralized approach as economies of scale are attained more easily. The study provides important indications of the factors at play in the SKU assignment problem, without however suggesting a detailed model that can be used in practical assignment problems.

In their 2018 article Holzapfel et al., propose an approach to solving the problem, presenting a case study of its application at a large German retailer. The supply chain network studied comprised a mix of central, regional and local warehouses. Each product is only assigned to one type of warehouse. Central warehouses replenish a large group of stores while regional and local ones only supply a smaller subset of stores, in a way that a product can, at any given time, be supplied by only a single warehouse. The authors start by noting that these decision are of tactical nature, belonging to a four to twelve months planning horizon, as altering product flows requires substantial physical and administrative tasks. They also focus heavily on the overarching implication that different options of product allocation have on costs across the supply chain, which calls for holistic approaches, that take all of these impacts into account. The proposed approach starts by identifying the subsystems of the supply chain affected by the SKU allocation, and subsequently finding for each of them the relevant cost drivers that need to be modeled. The groups of costs identified by Holzapfel et al. (2018) are inbound transportation costs, warehouse inventory costs, warehouse picking costs, outbound transportation costs and in-store operational costs. In the analysis preformed, cost drivers are classified as being dependent of the assignment of individual

SKUs or as resulting from the assignment of SKUs as a group.

The authors of the study further comment on the specific impacts of SKU assignment on the supply chain subsystems. Inbound transportation costs are influenced by the increased number of journeys needed to supply a larger number of smaller regional or local warehouses, which is a consequence of the arborescent structure supply-chain network studied. Warehousing costs are affected through different inventory needs and operation costs in the warehouses. It is referred that central warehouses replenish a larger number of stores, resulting in bigger pool of demand variations, which in turn affects inventory positions, cycle stocks and safety stock levels. This ultimately impacts the capital tied-up in system inventory. On the other hand, the cost of warehouse operations, namely picking tasks, varies from location to locations due to varying labour costs or technological constrains. Outbound transportation costs are influenced by different shipment quantities departing from each node of the supply chain, depending on the assignment preformed. It is also stated that in-store operations, a major cost driver in retail operations, might be influenced by the product mix of the received pallets, which is a direct consequence of the product range available in each warehouse. Shelf re-positioning benefits from pallets containing SKUs that are grouped by their position in the store's layout. In DCs that have a wider rage of products stocked, this can be more optimized.

3.4 The Backroom Effect

In the specific supply-chain structure studied, the SKU allocation impacts the amount of stock held at the backroom of the stores. It is therefore important to understand the cost implications of the Backroom effect.

Eroglu et al. (2013) explores the relation between orders arriving at stores, shelf space availability, the need of the store to keep merchandise in the backroom warehouse and the cost impacts that this brings.

According to Eroglu et al. (2013), backroom storage happens when there's relatively small allocated shelf space for a product and the large incoming orders make on-hold inventory surpass this allotted space. Larger average inventory levels are a direct consequence of order lead-times and order frequency, increasing with the increase of the first and reducing with increase of the latter. Every time there is overflow (defined as the inventory than surpasses the shelf's space) inventory is effectively split between two locations. As such, between replenishments, store management both have to check when shelf space gets released and the levels of backroom inventories and coordinate shelve refills from the backroom as purchases happen. Additionally, items can be misplaced or completely forgotten in the backroom, originating errors in the inventory records (which impacts the re-ordering algorithms) and lost sales, due to stock being in store but not actually in the shelf (creating stock-outs, with stock in store). It's argued that storing inventory in the backroom deeply changes the nature of the optimization of order quantities, which Eroglu et al. (2013) name as the "Backroom Effect" (BRE).

Eroglu et al. (2013) model in-store costs both leaving aside and considering the impact of the backroom effect in the costs and successfully prove the large cost gap of taking BRE in consideration when optimizing reordering quantities and policies.

3.5 Inventory Replenishment Policy

A large portion of the proposed model is based on the workings of the inventory policy used by the retailer. The inventory policy is a set of decisions regarding the amount and time of replenishment orders that a retailer makes and in which warehouses and stores inventory should be held. More specifically, multi-period inventory systems are designed to guarantee that an item will be available at the store when needed. The product will be ordered multiple times throughout the year, with the rules defined in the ordering system dictating the quantities ordered and the timing of the orders. (Jacobs and Chase, 2012)

Firstly, it's important to understand the concept of inventory position, which translates to the sum of stocks on-hand and of all outstanding orders that have still not arrived.

For the case study in focus, the replenishment relationship between stores and company's warehouse is the most relevant. The most common stock replenishment policies can be characterised by the allowed moments of re-order and by the quantities that are reordered. Regarding reordering moments, policies can be of (i) continuous review *(S)*, meaning that the inventory status is monitored continuously and orders are placed when inventory drops under a predefined threshold, or of (ii) periodic review *(R)*, meaning stock is reviewed and orders are placed only at specific, periodic intervals. Considering reorder quantities, systems can use a (i) fixed-quantity policy (Q), which only allows for reordering of a fixed predetermined quantity, or an (ii) up-to-level policy, which generates orders that take inventory to the a certain level, taking into account the inventory position at the reordering moment. (Jacobs and Chase (2012) , Silver et al. (2016))

One of the most widely used policies (including in the case study approached here) is the periodic review order-up-to-level policy (R,S). According to Silver et al. (2016), the control procedure is that at every review moment (e.i at every R time units) enough quantity is ordered to elevate the inventory position to the level S. For now let's assume that R and S are fixed values through time.

Silver et al. (2016) refer that often the value of R is determined by external factors, which happens to be the case in the supply-chain in study. The relevant question is how to effectively establish the value of S. At every review moment the inventory position needs to be elevated to protect sales for the period $(R+L)$, meaning equal to the review period and the order lead time. This is illustrated in Figure 3.2, adapted from Silver et al. (2016) where two consecutive orders are placed at moments t_0 and t_0+R, respectively, and arriving at the store at t_0+L and t_0+R+L, respectively. Let the orders be referred as X and Y. In moment t_0, as the S level is being selected, it must be recognized that, once order X has been processed, no other later orders (namely Y) can be received before t_0+R+L. As such, S at time t_0 must be enough to respond to demand through a period of duration $R+L$. Silver et al. (2016) also refers that a stock-out could happen in the early portion of the period (prior to t_0+L, i.e. the arrival of order X) but that would relate to the setting

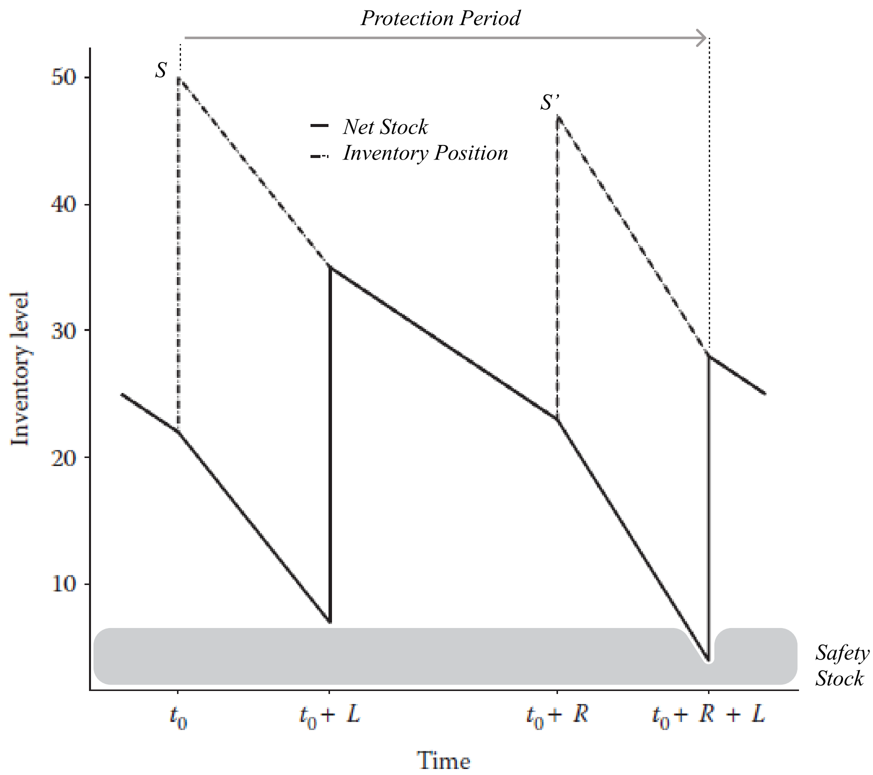

Figure 3.2: *(R,s)* policy exemplified, (Silver et al. (2016), *adapted*)

of the S value on the order preceding order X. The main point is that a stock-out happens toward the end of the period (after the arrival of order X) if the total demand in the period equal to $R+L$ is more than S.

A problem that comes up when setting up the system is that when placing orders demand is calculated using sales forecasts, which have inherent errors. The solution is to hold some extra stock, know as Safety Stock (*SS*), which must be included in *S*. In this way *S* is comprised of *SS* and expected demand, also know as Cycle Stock (*CS*). King (2011) notes that safety stock levels are calculated based on the expected distribution of forecast errors and order lead time, and are dependent of the defined service level, which sets the needed additional stock based on the intended probability of stock-outs, which in turn means that a certain number of stock-outs is naturally expected to happen, in extreme cases when demand is unusually high. For example, when setting a service level of 95% it's expected that in 5% of the order cycles demand will surpass the forecast Cycle Stock and Safety Stock. However, it's foreseeable that at least in 50% of cycles the Safety Stock remains untouched and that for the remaining 45% of cases the *SS* will be able to absorb the extra demand, and then will be compensated in the next demand cycle. As the inventory policy in study is a (R,S) policy, the calculation of the needed *SS* must also take into account the protection period of $(R+L)$, and should consider forecast errors for this time interval.

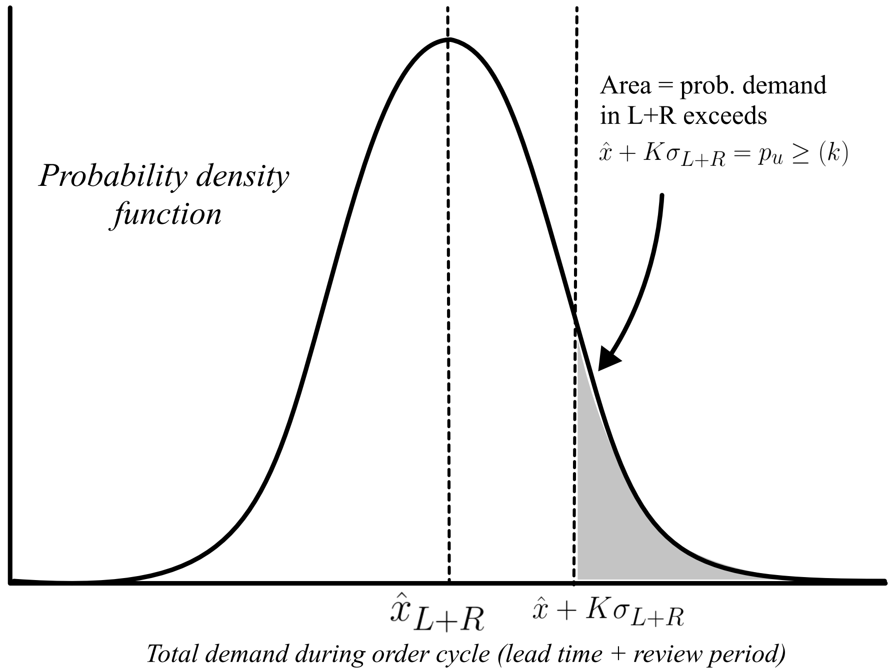

Figure 3.3: Curve of lost demand, (Silver et al. (2016),*adapted*)

(Silver et al., 2016) The *SS* is given by:

$$SS = F^{-1}(SL) \times \sigma_{R+L} \tag{3.1}$$

Here *SL* stands for the service level established for the product at a specific location and σ_{R+L} is the standard deviation of forecast errors for the period equal to $R+L$. Now let σ represent the standard deviation of forecast errors of a single period ahead of the order placement, consider that daily forecasts errors follow a normal distribution and that errors in consecutive days are assumed to be uncorrelated. If demand is assumed to be stationary, meaning the average demand is constant over time, the safety stock becomes:

$$SS = F^{-1}(SL) \times \sigma\sqrt{R+L} \tag{3.2}$$

Silver et al. (2016) follows the calculations of safety stock with an important derivation, which is the expected amount of lost sales. This is a direct consequence of the established value of *S*, which in turn depends on the amount of *SS* defined by the retailer. A stock-out will happen at a given replenishment cycle if actual sales for that cycle are higher than *S*, which corresponds to a probability given by the shaded area of Figure 3.3 and can be computed by equation 3.3.

$$Pr(Stock-out) = Pr(x \geq s) = \int_s^\infty f_x(x_0)dx_0 \tag{3.3}$$

While the chosen service level is set on the metric of stock-out probability (being considered an α type service level), the true cost of lost sales is better translated by the expected lost sales for a certain stock-out probability (i.e. the probability that a stock-out will happen at any given replenishment cycle multiplied by the expected loss of sales for that probability). This metric is known as Expected Shortage Per Replenishment Cycle (*ESPRC*) This can be expressed with Equation 3.4 that in turn equals the $Loss function$ of the normal distribution. Here k denotes a certain safety factor resulting from a defined service level as given by the inverse normal cumulative distribution in a way that $k = F^{-1}(SL)$.

$$ESPRC = \int_{S}^{\infty} (x_0 - S) f_x(x_0) dx_0 = \sigma_{R+L} G_u k \tag{3.4}$$

Another particularly useful observation made by Silver et al. (2016) is presented in Equation 3.5, where *DR* stands for demand during the review period.

$$E(OHinv) \approx SS + DR/2 \tag{3.5}$$

This happens because if the mean rate of demand is set at a fixed value, on average stock on-hand levels drop linearly form the moment just after a replenishment arrives, when they equal $SS + Avg(Order)$ to the moment right before the next replenishment arrives, when it reaches the value of *SS*. This equation will later be used in determining average on-hand inventories in stores.

As referred by Silver et al. (2016) the (R, S) inventory policy allows retailers to adapt to changing expected sales, at every R units of time, by adjusting the order-up-to-level S. This is a useful property if the demand pattern changes with time. In these cases the needed parameters come from the forecasting process from which we get the point forecast $x_{t,t+L+R}$ and standard deviation of forecast errors $\sigma_{t,t+L+R}$ for demand over the next $L + R$ periods. This means that it is possible for both Cycle Stock and SS to increase or decrease at each revision period. This is exactly how the retailer in study manages it's inventories, which brings another layer of complexity to the standard (R, S) policy.

3.6 Knapsack problem

After the supply chain cost modelling phase of the proposed solution, the problem of deciding which products to hold at the constrained warehouse takes the form of a classical optimization problem: the Knapsack problem, for which vast literature is available. On their reference work on the matter, Martello and Toth (1990) explore the problem on its multiple variants and gather possible ways to tackle them. The most generic definition of a knapsack problem pictures a limited container, a knapsack for example, where a series of objects need to be inserted. These objects both have a *weight* (i.e. they use a portion of the capacity of the container) and provide a certain *comfort* or *reward*(i.e. some form of profit variable that can be computed in an objective function). In the standard formulation of the problem, objects are numbered from 1 to n and are gathered in a vector of binary variables $X_i (i = 1, 2, ..., n)$ meaning the following:

$$x_i = \begin{cases} 1 \text{ , if the oject i is selected} \\ 0 \text{ , otherwise} \end{cases} \tag{3.6}$$

In this way, if p_i is the measure of profit given by object i, w_i its weight and c the size of the knapsack, the problem narrows down to selecting from i the items that satisfy the capacity constraint and maximize the objective function:

$$\sum_{n}^{i=1} w_i x_i \leq c \tag{3.7}$$

$$\sum_{n}^{i=1} p_i x_i \tag{3.8}$$

The problem has many different variations, depending on its exact structure. Variants include the admission of non-integer quantities for the objects, known as Multiple Choice Knapsack problems, or the use of more than one weight constraint, known as Multidimensional Knapsack problems (Fréville, 2004).

As observed by Martello and Toth (1990) knapsack problems are NP-hard. The authors go on to propose a series of approximate algorithms to determine feasible solutions, that can serve as lower bounds on the optimal solution.

One of the very first approaches suggested is a greedy heuristic based on the linear programming relaxation, that works by ordering items on decreasing value of profit per unit weight, so that:

$$\frac{p_1}{w_1} \geq \frac{p_2}{w_2} \geq \ldots \geq \frac{p_n}{w_n} \tag{3.9}$$

And then consecutively adding them to the container until the first item does not fit. The authors observe that the generated approximate solution can be quite close to the optimum solution. However, worst-case performance ration can be arbitrarily poor, depending on the instance. This method should be used with due caution, and bear in mind the specific instance in analysis.

Chapter 4

Methodology

The following chapter addresses the approach taken for the SKU allocation problem. Section 4.1 consists of an overview of the physical process that the algorithm was developed to support and the definition of the thesis' scope. After that, the chapter is divided along the two phases of the approach. In section 4.2, a cost modeling basis is described that is subsequently used to capture supply chain costs. Subsection 4.2.1 lists the model's assumptions and notations. Subsection 4.2.2 is concerned with the modelling of logistical, inventory, stock-out and in-store operations costs respectively. Section 4.3 describes the second phase of the approach, which is the greedy algorithm that uses the outputs of the cost modelling phase to suggest appropriate product allocations in the regional warehouse.

4.1 Allocation Process Overview

The proposed tool lays at the centre of the SKU allocation process of the retailer. This is the chain of tasks that happen from the moment the decision to change the allocation of a specific SKU until the moment the change is fully implemented. Naturally, the tool needs to work in accordance to the functioning of this higher level process.

Traditionally, the allocations have happened twice a year, making use of the previous year's sales records of each item to derive decisions.

For the new version of the allocation process the aim was set to make it faster, in a way that it could happen more often during the course of a year, thus better accompanying changing costumer demands. The main inputs for the centralizing algorithm were defined as the sales forecasts for the target centralization period.

The cost model and the mock-up produced were set to be independent from the TP, which equates to the target period in which is expected that the allocation will remain unaltered and the FAP, the flow alteration period which occurs between the allocation decision and the beginning of the target period, and is used by the retailer to enact the identified flow changes (see Figure 4.1) . This option was taken because the retailer intends to shorten both these periods in the future, in order o better adapt to shifting consumer demand and product seasonality.

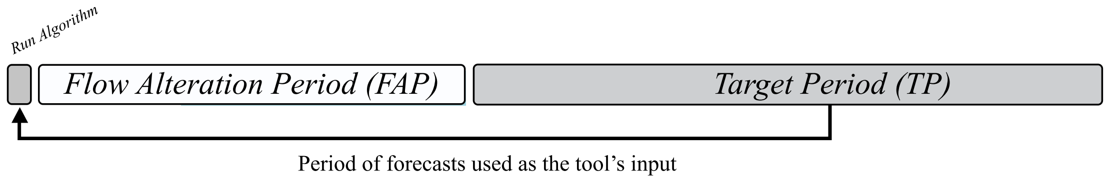

Figure 4.1: Scheme of expected centralization periods

The SKU allocation process includes tasks like changing replenishment inputs in the retailer's information systems, releasing storage positions in the regional warehouse by pushing the extra inventory to the stores in the region for the products that are being decentralized or sending a large first shipment of an item that is to start being stocked in the regional warehouse. The challenge then is how to best make the SKU allocation decisions to minimize overall supply chain cost.

The approach used in the case study was a combination of a supply chain model for each product followed by a simple algorithm to optimize the SKU mix of the regional warehouse (see Figure 4.2).

The first phase, corresponding to the cost model, focuses on extrapolating values for the cost and for the utilization rates of the constrained nodes of the supply chain. Note that this was done for both flows: the one that the product was historically assigned to and the alternative flow. The model includes realistic cost drivers incurred by the retailer for each of the two flow options available. The proposed methodology for SKU allocation is holistic, following a bottom-up approach.

Pragmatic decisions had to be made both on what to include and exclude from the model and how to create the best proxy for each included element. Only supply chain costs that are directly influenced by the product's allocation were included in the model. This means that costs that would occur independently of a product's allocation are not accounted for. Among these are for example warehouse and store fixed overheads, costs of receiving incoming inventory in the central warehouse or the cost of in store shelf replenishment (only the extra work needed, because products don't fit in the shelf is accounted for). As a result of this, when referring to Total Supply Chain Cost ($TotalSCC$) this really means the sum of costs that are directly derived from one allocation or another.

The methodology models, for each item i, in each store s, in each flow f, the supply chains costs resulting from warehouse manipulation tasks, transportation, inventory-holding (in stores and in the regional warehouse) and extra handling costs in stores due to stock overflows. It also models expected stock-outs caused by no available stock in the store. It purposefully leaves out the so called "stock-outs with stock" , that happen when there is stock available but the staff fails to place it in the shelves, and shrinkage costs. Reasons for this are further explained in chapter 7.

The final outputs of the cost modelling phase are the predicted $TotalSCC$ for the centralized flow and for the non-centralized flow, for each product i during the TP, and the use of storage positions in the regional warehouse.

The model used was tested against the historical data available, meaning compared with the flow historically used. After thorough verification of the assumptions , an initial unconstrained decision could be reached, which recommended an ideal individual allocation, but didn't account for the limits of the regional warehouse. On the second phase a simple algorithm that traded-off potential saving and capacity usage was used to achieve the best real world "solution", that as Holzapfel et al. (2018) notes, might differ very much from the unconstrained one.

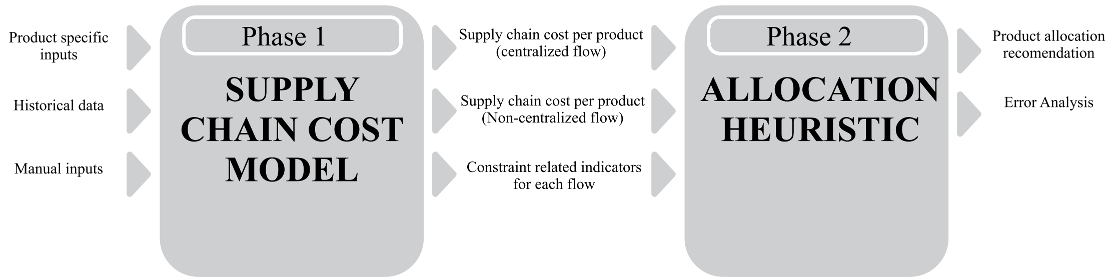

Figure 4.2: Two phases of the optimization process

4.2 Phase I - Supply Chain Cost Modelling

The explanation of the cost model will touch on the following topics in order: assumptions and simplifications made, model notation, calculation of some important variables used in several cost proxies and finally a detailed look into each individual cost proxy.

4.2.1 Assumptions and Model Set-up

The developed cost model relies on the following assumptions.

1. Order lead times are constant, which in practice does not happen.

2. If an item is assigned to the regional warehouse, it cannot be simultaneously supplied directly to the stores by the mainland warehouse.

3. It is assumed that during TP, daily demands follow a uniform normal distribution, centered around the average daily demand of the forecast period, which allows for the estimation of mean store inventories and consequently inventory costs.

4. It is assumed that during TP, daily forecast errors follow an uniform normal distribution and that errors in consecutive days are uncorrelated. As demand is assumed to be stationary, this in turn translates to daily forecast errors remaining unchanged. This allows for the estimation of the SS quantity.

5. No manual order overrides are preformed during the period: all orders are the direct result of the inventory policy placing orders in the information system.

6. Other important input elements that effectively change the order quantities, such as order quantity caps or the size of the Store-pack [1] are ignored.

7. All products assigned to the regional warehouse use a single picking position.

Expected inventory in stores

For determining inventory, stock-out and extra replenishment costs and to preform the optimization algorithm , it is important to analyse the store re-ordering process more in depth, with the goal of extrapolating values for the expected average on-hold inventory in each store of the region. This metric serves as the basis for the cost drivers of all these cost categories. Here only the flow of inventory between warehouses and stores is being consider, leaving the flow of goods between the supplier and the central warehouse out of the problem's scope.

The inventory policy used by the retailer, as said elsewhere in this text, is a variation of a (R,S) stocking policy, where the value of S is dynamically adapted at each review period to reflect the forecast sales.

Let's focus on the way the value of S is calculated, which from now on, according to notation will be defined as $Needs_{i,s,t}$, with the $Needs_{i,s,t}$ being dynamically adapted with sales forecasts to respond to forecast demand. The system works in the following manner: at every R time units the inventory position is checked and an order is placed so that IP is raises up the $Needs_{i,s,t}$ value. The necessities parcel comprises both stock held on stores as SS and expected sales for the protection period $(R+L)$, for reasons explained in chapter 3 (Silver et al., 2016). An additional element is added to the calculation of $Needs_{i,s,t}$, which differs from the traditional (R,S) policy. This is the $MinimumDisplayQuantity$ (MDQ), which is the minimum number of units needed in the product shelf to maintain good presentation to costumers. The formula effectively works by selecting the highest of the two possible values, MDQ or SS, and adding them to the expected Cycle Stock.

Let's consider the calculation of $Needs$, for one SKU i, in one store s in a review period. Equation (4.1) indicates the necessities and therefore the order quantity in a review period.

$$Q = Needs = max\{max\{MDQ,SS\} + CycleStock - IP\}, 0\} \tag{4.1}$$

(subscripts dropped for simplicity)

Both safety stock and cycle stock are functions of expected demand. These values effectively change with every replenishment period, which makes it hard to extrapolate the on-hand inventory at each moment without taking a unnecessarily complicated recursive approach. As the main objective is the determination of average on-hand inventories a simplification is used, where the expected value of daily demand is said to be constant and equal to the mean value of demand during the target period, TP. This effectively approximates the inventory policy to a standard

[1] Store-pack refers to the number of items contained in a box delivered to the store, functioning effectively as the Minimum Common Denominator for all orders. This amount is predefined and is the result of negotiations between the retailer and its suppliers

Table 4.1: Table of Notation

Indices	
i	SKUs sold in the Region, supplied from the mainland
s	Stores in the region
f	Flows that a product can be assigned to. Takes values $Cent$, for the centralized flow and $NCent$
Decision Variable	
x_i	For each product i, if x_i equals 1 the product flows a centralized flow, with stock being held at the regional warehouse. Otherwise, the product follows the Non-centralized flow. x_i can take value of 1 or 0.
Parameters	
Cost Group Parameters	
$SCC_{i,f}$	Total supply chain cost of supplying product i to the region, utilizing flow f during the target period.
$WhCost_{i,f}$	Total warehouse operations cost of supplying product i to the region, utilizing flow f during the target period.
$TCost_{i,f}$	Total transportation cost of supplying product i to the region, utilizing flow f during the target period.
$InvCost_{i,f}$	Total inventory cost of supplying product i to the region, utilizing flow f during the target period. Combines inventory costs at stores and at the regional warehouse.
$StockOutCost_{i,f}$	Total stock-out cost of supplying product i to the region, utilizing flow f during the target period.
$ReplCost_{i,f}$	Total extra handling cost associated with in-store shelf replenishment of supplying product i to the region, utilizing flow f during the target period.
Policy parameters	
$Q_{i,s,f}$	Average order quantity of product i to store s during the target period, when flow f is used for replenishment
$SS_{i,s,f}$	Average safety stock held of product i at store s during the target period ,when flow f is used for replenishment
$OHinv_{i,s,f}$	Average inventory of product i at store s during the target period ,when flow f is used for replenishment
$ESPRC_{i,s,f}$	Expected loss of sales resulting from zero stock of product i, at store s, when flow f is used for replenishment
$Overflow_{i,s,f}$	Expected excess of product units, when compared with the store PS at the moment of order arrival at the store,product i, at store s, when flow f is used for replenishment (in units)
$SSWh_i$	Average safety stock held of product i at the regional warehouse.
$OHinvWh_i$	Average inventory of product i held at the regional warehouse

Table 4.2: Table of Notation - (cont.)

Period Parameters	
TP	Duration of target period in days, excluding days when stores are closed
L_f	Replenishment lead time, from warehouse to store in flow f
R_f	Replenishment review period of stores, in flow f
L_{Wh}	Replenishment lead time from mainland warehouse to regional warehouse.
R_{Wh}	Replenishment review period of regional warehouse.
Forecast Parameters	
$D_{i,s}$	Average daily demand forecast for product i in store s during the target period, in units
$ForErr_{i,s}$	Coefficient of variation of forecast errors per product i per store s
$SL_{i,s}$	α Service level set by the retailer per product i in store s
$SLWh_{i,s}$	α Service level set by the retailer per product i for replenishment of the regional warehouse by the central warehouse.
Logistical Operation Parameters	
$PS_{i,s}$	Presentation stock of product i in store s set by the retailer (in units). Corresponds to maximum allocated shelf capacity.
$MDQ_{i,s}$	Minimum display quantity of product i in store s set by the retailer (in units)
$Storepack_i$	Nr of units of product i contained inside of a box shipped to stores (in units)
$Boxvol_i$	Physical volume of a box of products i, in m^3
$Palletcap_i$	Nr of boxes of a supplier pallet for product i, in number of boxes
$PalVolume$	Total volume that a PBS prepared pallet can carry in, m^3
$Palletutil$	Average percentage of real volume used in a PBS prepared pallet
$Containerutil_f$	Average number of pallets shipped container in the boat per flow type
$Pick_{i,DC}$	Cost of PBS picking product i at a specific warehouse. DC can equal $Central$ or $Regional$.
$Xdocking$	Cost of crossdocking operation per pallet in the regional warehouse
$Dispatch$	Cost of dispatching a single product pallet in the Central warehouse
$Provisioning$	Cost of provisioning a single product pallet in the Regional warehouse
C_i	Unit cost of SKU i
P_i	Avg selling price of SKU i in the region's stores
$IHCS$	Inventory Carrying cost of stores for the period, in percentage of value of inventory (%)
$IHCDC$	Inventory Carrying cost of Warehouses for the period, in percentage of value of inventory (%)
$EHCB$	Extra-handling cost per box that corresponds to overflow. [a]
$Boattransp$	Cost of shipping one full container to the region, in
$Trucktransp$	Cost of transportation of one pallet from the regional warehouse to the regions stores, in
$LostSalesPct$	Percentage of lost demand due to stock-outs that is accounted for as actual lost sales (%)

[a]This cost accounts for all handling operations associated with boxes that surpass the PS value. These operations include for example trips to the backroom, box sorting, extra shelf replenishment.

(R,S), allowing for use of some known derivations of this policy, as listed in chapter 3. Values of expected inventory still change with both flows,as these have different order lead-times and replenishment windows, but are now based on a constant expected demand.

It therefore follows from equation 3.3 now including the MDQ factor:

$$E(OHinv_{i,s,f}) \approx max(SS_{i,s,f}; MDQ_{i,s}) + DR_{i,s,f}/2 \tag{4.2}$$

For the calculation of SS it should be noted that the company keeps records of average coefficient of variation of daily forecast errors for each i product in each s store . The standard deviation of errors for an order, σ_{L+R}, will equate to the sum of the square roots of the daily standard deviations of the period. As it was defined that demand and forecast errors were to be constant through time, the expression for expected standard deviation of forecast error for one SKU in a store can be developed into:

$$\sigma_{L+R} = D \times ForErr \times \sqrt{L+R} \tag{4.3}$$

(subscripts dropped for simplicity)

The retailer then computes the appropriate SS, by calculating the additional stock needed to prevent a stock-out in a day with a defined probability equaling the Service Level (SL). By its characteristics this is an α type service level, as it protects against the occurrence of stock-outs and not depth of lost sales. k denotes a certain safety factor resulting from a defined service level as given by the inverse normal cumulative distribution in a way that $k = F^{-1}(SL)$. The additional stock is given by:

$$SS = \sigma_{L+R}k \tag{4.4}$$

(subscripts dropped for simplicity)

The only element from Equation 4.2 that still needs to be calculated is the demand during the review period. As daily demand is defined as constant, the demand during the review period can be calculated by multiplying expected daily demand by the number of days in the review period, as demonstrated in Equation (3.5) Finally, joining the previous elements we have that the expected On-hand Inventory of SKU i, in a store s, in flow f, can be calculated by:

$$E(OHinv) \approx max(\sigma_{L+R}k; MDQ) + (D \times R)/2 \tag{4.5}$$

(subscripts dropped for simplicity)

The points mentions above are graphically represented in Figure 4.3. The next steps consist in determining values of average order and safety stocks for each flow, using as input expected sales and the characteristics of each replenishment flow, such as order lead time and order review period.

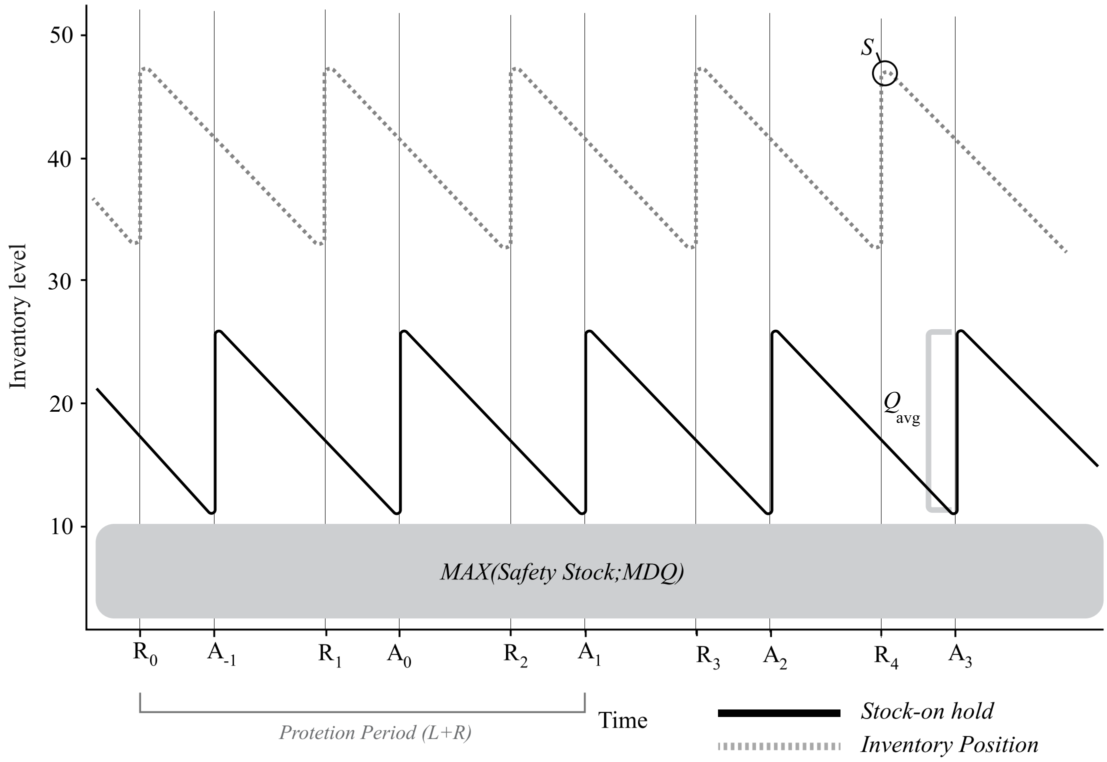

Figure 4.3: Illustrative example of the simplified inventory policy modeled

Expected inventory at the regional warehouse

The expected stock levels in the regional warehouse are obtained in a similar way to the stock estimations for stores, as will be shown the following paragraphs.

It's defined that the regional warehouse should provide stock to the stores with a certain α service level, translating the probability of a stock-out during a replenishment cycle of the warehouse.

We proceed in determining what the distribution of demand in the regional warehouse will be, for a single day of operation, for a product i. The daily demand distribution will equate to the combination of demand distributions of individual stores, which are assumed to be constant and independent. It's also assumed that stores are replenished every day, so that demand at stores is "pulled" from the warehouse.

The resulting distribution will have an expected value equal to the sum of sales forecasts for individual stores and a variation of forecasting errors that will be equal to the sum of variations of forecasting errors for individual stores. Note that expected sales values are given by the expected sales forecast and that the standard deviation of forecast errors are extracted from the historical coefficient of variation of forecast errors. The daily demand in the region for a product i then becomes the daily demand at the warehouse and is assumed to follow a normal distribution with:

$$\mu_{DailyRegionalDemand} = \sum_{s}^{stores} D_{i,s} \tag{4.6}$$

$$\sigma_{DailyForErrRegDem} = \sqrt{\sum_{s}^{stores} (D_{i,s} * ForErr_{i,s})^2} \tag{4.7}$$

Then, applying the same principle of equation 3.5, the goal is to find the values of safety stock and demand in the regional warehouse over the review period. Safety stock needs to protect against forecast errors during the protection period of $R_{Wh} + L_{Wh}$, that is shorter for the warehouse than for the stores. Assuming constant sales, independent and constant forecast errors between days of the protection period the standard deviation of errors can be computed as being:

$$\sigma_{RPForErrRegDem} = \sqrt{R_{Wh} + L_{Wh}} * \sigma_{DailyForErrRegDem} \tag{4.8}$$

Safety stock is necessary for the protection against large forecast errors. It is computed by multiplying $\sigma_{RPForErrRegDem}$ by a factor k, given by $k = F^{-1}(SL_{Wh})$.

To better model the situation at the warehouse, it was defined that the lower bound for safety stock of any product would be equal to the capacity of a full pallet, which translates to the permanent existence of at least one reserve pallet and one picking pallet. Regional warehouse SS is therefore computed by:

$$SSWh_i = max(\sigma_{RPForErrRegDem} * k; Palcap_i * Storepack_i) \tag{4.9}$$

Following the similarities with (3.5) the needed cycle stock is then given by half of the expected quantity order in the replenishment window, which is given by the daily demand multiplied by R_{Wh}. It follows that average inventory at the regional warehouse equates to:

$$E(OHInvWh_i) = max(\sigma_{RPForErrRegDem} * k; Palcap_i * Storepack_i) + R \times \mu_{DailyRegionalDemand} \tag{4.10}$$

4.2.2 Model's Cost-Proxies

Warehousing and Transportation Costs

Warehouse and transport operations deal with the movement and assortment of products inside the DCs and shipments between DCs and stores respectively, which translates to fairly similar cost modelling strategies for these cost groups. These movements are done using individual units, boxes, full pallets or full truck and shipping containers loads.

Because different allocations result in different chains of operations it was decided that the basis for cost drivers in each flow would be determining the number of box, pallet or container ship equivalents for the expected sales of a SKU i in the target period and then multiplying by the base cost of each operation. The costs associated with the chain of operations of would then be summed to obtain the total warehousing and transportation costs for each specific flow.

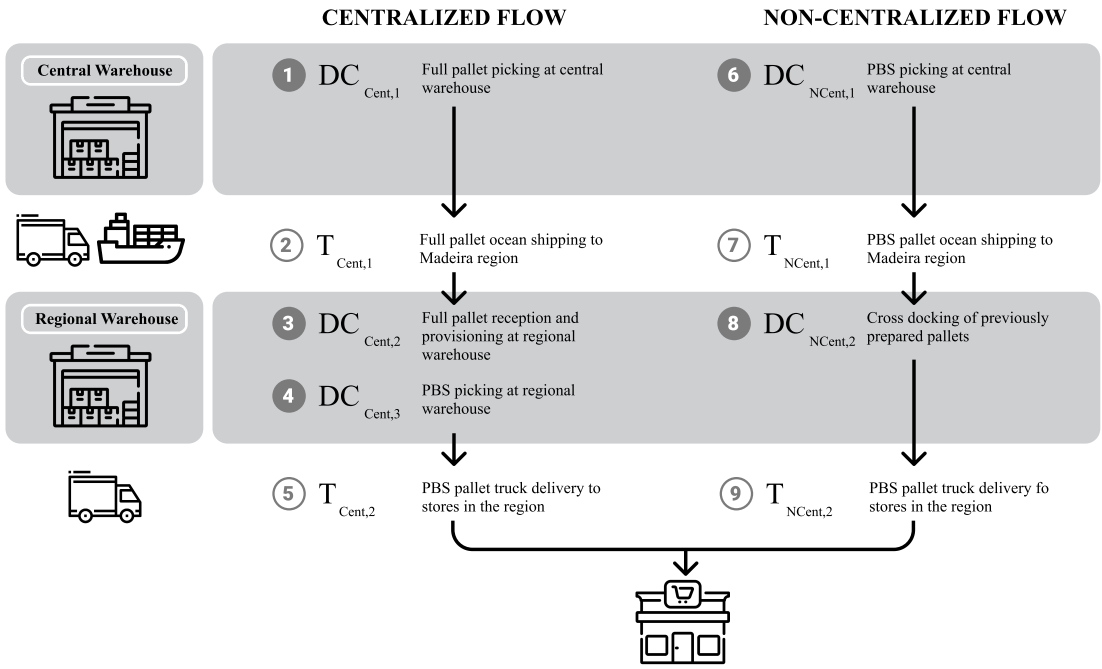

Figure 4.4: Scheme of transportation and warehousing tasks

For some operations these inputs are not that direct and some additional clarifications are needed. Costs of PBS picking per box are known at a category level, so the same cost per box is assumed to be the same for all SKUs belonging to the same category, though there are differences in costs between the regional and the main warehouse. For movements that use pallets, cost modelling becomes hard when dealing with PBS pallets, as these carry a mix of many different products. As such a proxy is reached by determining the volume of boxes in cubic meters for each SKU and dividing the value by the volume capacity of a pallet. If preformed for all products, the overall quantity of pallets transported predicted by the model is expected to be close to the number of PBS pallets actually transported. Stacking different shaped and sized boxes in a PBS pallet ultimately leads to loss of useful volume, as spaces accumulate between boxes. The parameter $Palletutil$ is used to adjust the full available volume of the pallet to the actual useful volume. The parameter $Containerutil_f$ is used to distinguish the average number of pallets carried by container ship per flow. This factor is needed because single product pallets, shipped in the centralized flow often are stacked on top of each other, whereas PBS pallets from the non-centralized flow can not be stacked and are limited by the container's floor area. This again represents an abstraction, as commonly both types of pallets are shipped in the same shipping container.

Figure 4.4 sums-up all of the warehouse and transport operations whose costs are modeled. These are followed by the respective cost expressions that represent the costs of each single warehousing or transport operation that are needed for a specific product i, during the target period. Equations 4.11 to 4.14 aggregate all these cost together into a warehousing and a transport cost, per product per flow, during the target period.

Centralized Flow

Nr	Proxy name	Base unit	Cost expression
1.	$DC^{Cent,1}$	$Pallet$	$\frac{\sum_s D_{i,s} \times TP}{Storepack_i \times Palletcap_i} \times Dispatch$
2.	$T^{Cent,1}$	$Container$	$\frac{\sum_s D_{i,s} \times TP}{Storepack_i \times Palletcap_i \times Contrainerutil_{f=Cent}} \times Boattransp$
3.	$DC^{Cent,2}$	$Pallet$	$\frac{\sum_s D_{i,s} \times TP}{Storepack_i \times Palletcap_i} \times Provisioning$
4.	$DC^{Cent,3}$	Box	$\frac{\sum_s D_{i,s} \times TP}{Storepack_i} \times Pick_{i,DC=Regional}$
5.	$T^{Cent,2}$	$Pallet$	$\frac{\sum_s D_{i,s} \times TP}{Storepack_i} \times \frac{Boxvol_i}{Palletutil \times PalVolume} \times Trucktransp$

Non-Centralized Flow

Nr	Proxy name	Base unit	Cost expression
6.	$DC^{NCent,1}$	Box	$\sum_s D_{i,s} \times TPStorepack_i \times Pick_{i,DC=Central}$
7.	$T^{NCent,1}$	$Container$	$\frac{\sum_s D_{i,s} \times TP}{Storepack_i} \times \frac{Boxvol_i}{Palletutil \times PalVolume \times Containerutil_{f=Cent}} \times Boattransp$
8.	$DC^{NCent,2}$	$Pallet$	$\frac{\sum_s D_{i,s} \times TP}{Storepack_i} \times \frac{Boxvol_i}{Palletutil \times PalVolume} \times Xdocking$
9.	$T^{NCent,2}$	$Pallet$	$\frac{\sum_s D_{i,s} \times TP}{Storepack_i} \times \frac{Boxvol_i}{Palletutil \times PalVolume} \times Trucktransp$

In expressions 1 to 9 i equals to the specific product that is being analysed.

Total warehousing costs per product in each flow correspond to:

$$WhCost_{f=cent,i} = DC^{Cent,1} + DC^{Cent,2} + DC^{Cent,3} \tag{4.11}$$

$$WhCost_{f=Ncent,i} = DC^{NCent,1} + DC^{NCent,2} \tag{4.12}$$

Total transport costs per product in each flow correspond to:

$$TCost_{f=cent,i} = T^{Cent,1} + T^{Cent,2} \tag{4.13}$$

$$TCost_{f=Ncent,i} = T^{NCent,1} + T^{NCent,2} \tag{4.14}$$

Inventory cost

The retailer will incur in inventory costs through inventory held in the backroom of stores and through the inventory held in the regional warehouse, with these costs shifting depending on the chosen SKU allocation. It can be argued that using a non-centralized flow will increase needed inventories in the mainland warehouse, but this is left outside of the scope because the regular delivery patterns of suppliers to the mainland warehouse mean that these extra inventory needs could be better managed by later purchases or passed onto the suppliers.

Focusing first on the stores, total store inventory of item i for the region equals the sum of the average inventory per store s, computed using equation 4.5, for all of the region's stores. This is then multiplied by the inventory carrying cost, adjusted for the target period and the unit cost of SKU i.

$$StInvCost_{i,f} = \sum_{s=1}^{Regstores} E(OHinv_{i,s,f}) \times C_i \times IHCS \tag{4.15}$$

Inventory cost in the regional warehouse of item i is given by the average inventory held at the regional warehouse, which then is multiplied by the inventory carrying cost (adjusted for the target period) and the unit cost of SKU i.

$$WhInvCost_i = E(OHinvWh_i) \times C_i \times IHCDC \tag{4.16}$$

The final step is adding both elements together to obtain the total inventory cost. Warehouse inventory cost is only accounted for if the flow being used is the centralized.

$$InvCost_{i,f} \begin{cases} StInvCost_{i,f=NCent} \text{ If flow is Non-centralized} \\ StInvCost_{i,f=Cent} + WhInvCost_i \text{ If flow is Centralized} \end{cases} \tag{4.17}$$

Expected stock-outs

The retailer defines a certain level of safety stock to provide protection against days when actual sales deviate a lot from the forecast. It was also mentioned that this will generate a certain level of expected days when demand will surpass stock available in the store. Because of the different inventory policies of the flows, there will be an impact on this type of stock-outs.

In chapter 3, a way to quantify ESPRC for a given (R,S) policy was suggested. As with the equation for average in store inventory, the expression needs to be adapted for existence of minimum display quantities in the stores shelves.

Each cycle will have expected losses equalling $L(Z_s)$ where Z_s is the $Zscore$ corresponding to the order up-to-level value of S. Note that S will equal the maximum value between $CycleStock + MDQ$ and $CycleStock + SS$.

$$ESPRC_{i,s,f} = L(Z_{i,s,f}) \tag{4.18}$$

Total losses for the target period can be calculated by multiplying the ESPRC by the number of replenishment cycles inside of the target period.

As for the economical valuing of lost sales, their impact depends on the consumers behavior when confronted with a stock-out. Some consumers will postpone the purchase, some might buy substitute products while some will not purchase at all or buy at competitors, representing a true loss of sales. The loss only affects the profit margin of true lost sales and not the full revenue. Everything being considered, the cost of zero-stock stock-outs for product i, in store s, for flow f during the full target period is given by:

$$StockOutCost_{i,f} = (\sum_{s=1}^{Regstores} ESPRC_{i,s,f}) \times \frac{TP}{R} \times LostSalesPct \times (P_i - C_i) \tag{4.19}$$

In-Store Extra Handling Costs

As noted elsewhere in this document, in-store operations represent a big component of the costs incurred by the retailer with SKU allocation having an indirect impact on these costs through the average inventory on hand. As an order for a given product i arrives at store s, there is a process to place the received merchandise on the shelves. If when this happens the space available in the shelf is not enough to hold all units, extra tasks must take place, such as taking the extra inventory to the backroom, sorting the extra units in the backroom and repeating the shelf placement when enough purchases happen, freeing up the needed shelf space. A very exact modelling strategy is hard to quantify because it is directly dependent on the different ways replenishment operations are done in each store. Note that this is repeated for every replenishment cycle.

For the purposes of the model, it was decided to determine the expected overflow of boxes at the moment that an order is received and defining a fixed extra handling cost per box, that encompasses the cost of all extra handling tasks. At each replenishment cycle, for a specific combination of product i,store s and flow f the overflow value is computed by the difference between the expected on-hand inventory at the moment of order arrival and the value for the PS. If the value is positive it translates to overflow. If the value is negative there is no overflow and no extra cost. Note that the expected on-hand inventory is obtain by using equation 4.5, but using Q (i.e. the value for average order) instead of $Q/2$. Equation 4.14 translates to just this:

$$E(Overflow) \approx max(0; max(SS; MDQ) + (D \times R) - PS) \tag{4.20}$$

(subscripts dropped for simplicity)

The total extra handling cost for product i in flow f is then given by the sum of overflow boxes for every store in the region, multiplied by the extra handling cost per box and by the number of replenishment cycles in the target period.

$$ReplCost_{i,f} = \sum_{s=1}^{Regstores} E(Overflow_{i,s,f}) \times EHCB \times \frac{TP}{R} \tag{4.21}$$

This concludes the supply chain cost modeling phase. All proxies previously listed can be added together, resulting in the total supply chain cost per product per flow in the target period.

$$SSC_{f,i} = WhCost_{f,i} + TCost_{f,i} + InvCost_{i,f} + StockOutCost_{i,f} + ReplCost_{i,f} \tag{4.22}$$

4.3 Phase II - Prioritization Of SKU Allocation

After accurately modeling costs for each product i on the two flows, for a certain target period, the second phase of deciding which flow to use for each product bearing in mind the physical limitations of the regional warehouse can begin. In the next section the optimization phase of the problem is laid out.

Centralization savings can be calculated using the costs of both possible flows, as represented in Equation 4.23. The predicted usage of capacity in the regional warehouse, that is used to adjust allocation to the constraints is also an output of the first phase. In the generalization made for the problem products can make use of several picking positions, although in the case study in question this was simplified, assuming that $WhPick_i = 1$ for every product. For the reserve positions the calculation can be based on the predicted $OHinvWh_i$. To do this, the average inventory is rounded up to the next multiple of pallet capacity in units, and the number of used picking positions is subtracted. This is captured in Equation 4.24

$$CentSav_i = SSC_{f=Ncent,i} - SSC_{f=Cent,i} \tag{4.23}$$

$$WhRes_i = Roundup(OHinvWh_i/(Palletcap_i * Storepack_i)) - WhPick_i \tag{4.24}$$

At this point the problem transforms into a typical 1/0 knapsack problem, where the goal is to assign a set of items with a certain capacity use and a certain profit associated to them to a constrained space, as previously explained in chapter 3. In this case, the items to be assigned are the SKUs, the constrained space is the regional warehouse and the profits are the savings associated with assigning each individual item. The problem is actually a form of a multi-dimension knapsack problem as two variables act as *weight*, the use of picking positions and the use of reserve positions. The problem is expressed in Equations 4.25 to 4.28, and the used notation is clarified in Table 4.3.

Objective function

$$\mathbf{Min} z = \sum_{i=1}^{n} x_i \cdot CentSav_i \tag{4.25}$$

Constraints

$$\sum_{i=1}^{ProductRange} x_i \times WhPick_i \leq MaxPick \tag{4.26}$$

$$\sum_{i=1}^{ProductRange} A_i * WhRes_i \leq MaxReserves \tag{4.27}$$

$$\text{For all } i \; x_i \in 1;0 \tag{4.28}$$

Many possible algorithmic approaches are available in the literature to solve this type of problems in an imperfect but fast way, providing various degrees of approximation to the optimum solution, degrees of complexity and speed. Taking into account the input variables and constraints of the specific problem at hand, and for reasons further explain later, the used approach is based upon the simple greedy algorithm defined in chapter 3.

As decision variables only take values one or zero, it can be derived that the solution space of the problem is 2^n, with complexity scaling exponentially with the increase of product range. The approach is summed-up in table 4.2 and in Algorithms 1 and 2. After running the two greedy algorithms, each trying to approximate the optimum solution by ordering items for their specific constraint, the model chooses between the two provided allocations by selecting the one that produces the highest value of captured savings.

Assignment Change proposal

The very last step of the suggested allocation algorithm is presenting the proposed allocation solution as the list of needed alterations of the current state of product allocation, instead of the static proposal of ideal allocation. As such, it's indicated for each product i if their allocation should be maintained or changed, depending if the allocation suggested by the model differs from the current position.

As at every given moment many of the SKUs are already previously assigned to the flows that minimize supply chain costs, a great part of the modeled savings associated with product centralization are, in a sense, already "captured". A prediction of savings of any given centralization cycle is given by running the cost function for the real "as-is" product assignment at the moment and the product assignment proposed by the model.

Table 4.3: Table of Notation - Optimization phase

Decision Variable	
x_i	For each product i, if x_i equals 1 the product flows a centralized flow, with stock being held at the regional warehouse. Otherwise, the product follows the Non-centralized flow. x_i can take value of 1 or 0. Size of vector is n, which equals de number of products in study
Input Values	
$CentSav_i$	Difference in total supply-chain costs between the Non Centralized Flow and the Centralized Flow.
$WhPick_i$	Average usage of picking positions in the regional warehouse for a product i if the centralized flow is used.
$WhRes_i$	Average usage of reserve positions in the regional warehouse for a product i if the centralized flow is used.
$MaxPick$	Available picking positions in the regional warehouse.
$MaxReserve$	Limit of mean occupation of reserve positions in the regional warehouse .

Result: Greedy Reserve Allocation proposal
Model Inputs: n , $MaxReserve$, $MaxPick$, $WhRes_i$, $WhPick_i$, $CentSav_i$;
Model Outputs: x_i , $Savings$;
Begin ;
$PickCap := MaxPick$;
$ResCap := MaxReserve$;
$Savings := 0$;
sort i **by** $CentSav_i/WhRes_i$ **in descending order**;
for i = 1 **to** n **do**
    if $CentSav_i > 0$ **then**
        if $WhRes_i > ResCap$ ***and*** $WhPick_i > WhRes_i$ **then**
            $x_i = 1$;
            $Savings = Savings + CentSav_i$;
            $PickCap = PickCap - WhPick_i$;
            $ResCap = ResCap - WhRes_i$;
        else
            $x_i = 0$
        end
    end
end

Algorithm 1: Greedy Algorithm - Reserve restriction

Result: Greedy Picking Allocation proposal
Model Inputs: n , $MaxReserve$, $MaxPick$, $WhRes_i$, $WhPick_i$, $CentSav_i$;
Model Outputs: x_i , $Savings$;
Begin ;
$PickCap := MaxPick$;
$ResCap := MaxReserve$;
$Savings := 0$;
sort i **by** $CentSav_i$ **in descending order**;
for i = 1 **to** n **do**
    if $CentSav_i > 0$ **then**
        if $WhRes_i > ResCap$ ***and*** $WhPick_i > WhRes_i$ **then**
            $x_i = 1$;
            $Savings = Savings + CentSav_i$;
            $PickCap = PickCap - WhPick_i$;
            $ResCap = ResCap - WhRes_i$;
        else
            $x_i = 0$
        end
    end
end

Algorithm 2: Greedy Algorithm - Picking restriction

Chapter 5

Case Study at a Retail Company

The present chapter intends to present the relation between the proposed approach and the company environment in which it was developed. This is quite relevant as the approach must be intrinsically linked with the supply chain structure of the company. While a take on the SKU allocation task can be generalized for other situations, the built model needs to reflect the behaviors and relations in a very specific way. The chapter will follow this structure: in section 5.1, the company at large and the specific regional operations are described. The process that this thesis aims to improve - the SKU allocations to the regional warehouse - is also described in its current form.

Section 5.2 explores the architecture of the minimum viable product for the decision support system that implements the proposed approach and that produced the results explored in chapter 6. Finally, section 5.3 focuses the most important process inputs of both possible allocation flows that were used to build the model.

5.1 The Retailer's Operation Structure and As-Is Process

SONAE is a Portuguese based multinational company that runs a diversified portfolio of companies with various degrees of control and in diversified businesses with big focus given to areas like food and specialized retail, financial services, technology, shopping centers and tel-com. (SONAE, 2021c)

SONAE was initially founded as *Sociedade Nacional de Estratificados* in 1959, starting as a mainly industrial company. Throughout it's life, SONAE progressively expanded to other business sectors, ultimately becoming one of the largest Portuguese business groups. In the mid-1980 SONAE entered the retail sector by opening the country's first hypermarket, in Matosinhos, with a rapid growth in the number of stores ever since. (SONAE, 2021b) SONAE's retail division is called SONAE MC and it's the leader in the food retail market in Portugal, it operates a series of supermarkets and hypermarkets under the brands *Continente* (41 locations, hypermarket brand) , *Continente Modelo* (132 locations, supermarket brand) and *Continente Bom-dia* (98 locations and observing rapid growth in new locations, proximity supermarket brand), as well as a series of adjacent businesses of specialized retail and healthcare (SONAE, 2021a). To support its retail operation, SONAE MC uses four main distribution centers located in Maia, Azambuja, Carregado and Madeira, with some minor locations in Algarve and the Azores islands and meat and fish production centers that work independently.

The retailer's regional operation

The allocation problem focuses on SONAE MC's operation in the Madeira Archipelago, so further focus needs to be given to this part of the retailer's operation. At the beginning of 2021, the retailer operated 15 mid-size supermarkets in the Autonomous Region of Madeira, exclusively on Madeira Island, under the *Continente Modelo* Brand. To support the operation of these stores, *SONAE MC* established a warehousing operation at *Água de Pena*, on the East of the Island, from now on referred as *AP* . Figure 5.1. presents a scheme of the operation described above. *AP* is small when compared with warehouses in the mainland. At *AP*, 3 different separate warehouses exist, consisting of an ambient temperature warehouse, a controlled positive temperature warehouse focusing on fresh products and a frozen products warehouse. *AP* runs two types of order prepping operations across the different operation temperatures: separate PBL operations for ambient room products and for fresh products and separate PBS operations for ambient room products, for fresh products and for frozen products. It also runs cross-docking operations for pallets that are already sorted by store (both in PBL and PBS) and directly shipped from mainland Portugal.

Products sold by *SONAE MC* in Madeira are supplied in different manners. Local suppliers provide some SKUs in all three operations (ambient, fresh and frozen) delivering both to the stores directly and to the warehouse at *AP*. Nevertheless, the majority of SKUs is supplied from the mainland warehouses through maritime shipment, in deliveries that happen twice a week on Tuesdays and Fridays. *Azambuja* warehouse supplies fresh and ambient temperature products, while *Maia* warehouse is the supplier for frozen products. In the warehouses of the mainland, the orders are prepared for the archipelago stores using PBL or PBS flows, but prepared pallets are shipped together in the same shipping containers and when they arrive in *AP* they are redistributed to their destiny stores, in a cross docking operation. Additionally, for products that follow a PBS flow, the company has the option of sending stock in bulk to the AP warehouse and supply the stores through the PBS operation there, instead of preparing them in the mainland.

Deciding which products should be assigned to the Madeira warehouse and which should be directly replenished from the mainland is the focus of this thesis.

Although this decision needs to be made for all items supplied from the mainland using a PBS flow, the project scope is limited to ambient temperature items that respect the previously outlined conditions. The limited scope intends to simplify the development of the proposed tool, with plans to later expand the scope to the frozen product category.

For the case study preformed 2594 SKUs were considered as corresponding to the criteria set above, with 6 being excluded from the analysis due to lack of critical data. In table 5.1, the studied items are broken down between the ones that are currently centralized and non-centralized. Data is also provided regarding the physical volume that these SKUs represented, with values of

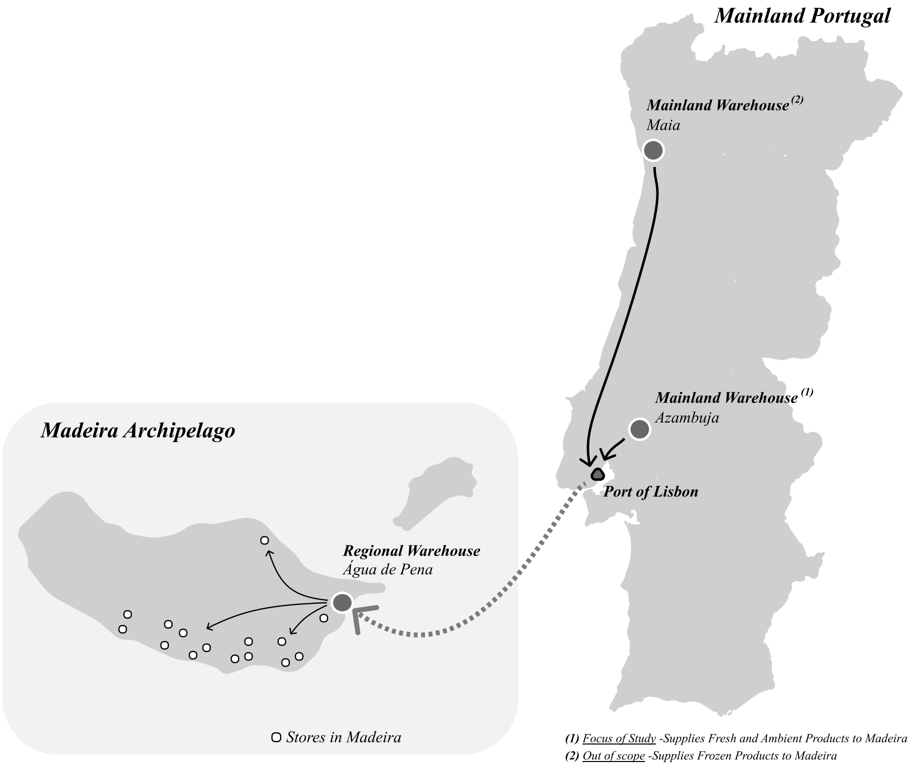

Figure 5.1: Operations in Madeira

boxes and pallets moved corresponding to sales in Madeira during analysed period (January and February of 2021)

The current process for SKU-warehouse allocation

The SKU-warehouse allocation is a task that comprises both the decision of where to allocate each item and the subsequent implementation of that change in the supply chain. Although this thesis focuses only on the decision algorithm component, it is important to know how the downstream process works.

Currently the process runs two times per year, in the beginning of summer and before the Christmas holidays, in order to reflect the big shifts in demand due to highly seasonal products. The process starts by running the current SKU allocation algorithm. It uses historical sales data to rank products according to total units sold and physical volume moved (in number of pallets processed) in the previous year. It then classifies products in A-B-C-D-F categories for each of the criteria mentioned above. Products that fall into categories representing more movement (A,B,C...) are assigned first to the regional warehouse first, until the cumulative number of predicted storage positions needed reaches the capacity of either the available picking positions or reserve positions.

The algorithm then compares the suggested placement of the items with their allocation at

Table 5.1: Overview of products inside of scope

		Centralized		Non centralized	
	Totals	**Absolute**	%	**Absolute**	%
NR of SKUs	2594	505	19%	2090	81%
NR of Boxes	466422	336141	72%	130281	28%
NR of Pallets	1933	1613	83%	319	17%

Initial Extraction	2594
SKUs removed due to no Sales during period	6

the moment, identifying which items are to be centralized[1], de-centralized[2] or kept at the same warehouse. These suggestions are then manually reviewed by the several process stakeholders involved in the process.

After the decision on which items to reallocate is made, a four week period follows. In the first two weeks, storage positions of products that will leave *AP* are released, by pushing extra stock to the stores and stopping bulk stock orders. In the following two weeks, products newly allocated to AP are subject to a large first delivery, in order to start PBS picking operations in the regional warehouse.

Motivation to start this project of creating a new allocation algorithm relates to some of the shortfalls that SONAE identifies in the current process. The main ones are the following:

1. Unusually large levels of stock-outs and store-backroom occupation observed in the stores in Madeira, when compared with mainland locations, which could potentially be improved by a better mix of products stored at *AP*.

2. The perception that the current method is too simplistic, by not accounting for important costs that result from the decision to allocate the products or not.

3. The general impracticality of the current process, that due to its poor design and need of manual revision, keeps the allocation process from happening more often during the year, which would better match changing consumer demands.

[1] Product previously held at mainland warehouse starts to be stored at the Regional warehouse
[2] Product previously held at regional warehouse stops being stored there

5.2 MVP for the Allocation Tool

One of the biggest challenges with the project was the definition of its exact goals and direction, representing a large portion of the time when compared to the actual development of the final tool proposed. The first phase was one of understanding how each of the studied product flows worked and how the process of SKU allocation change took place. The outputs of this investigation were process maps of the process and deep knowledge of the inner-works and shortfalls of the current process.

In the second phase, the effort was to define project goals, constrains and scope. This was difficult, as SKU allocation produces conflicting effects by nature across the supply chain, effectively shifting work and effort form a subsystem of the supply chain to another. Reflecting these trade-offs in an holistic, transparent and representative manner for all process steps and stakeholders was a real challenge. Simultaneously, extensive bench-marking with both the existing literature and with similar previous internal projects was done. In the end of this phase, a decision was made to model all relevant costs that each item would incur if it would be a assigned to one allocation or another, as well as the elements tat would act as constraints, such as average inventory levels at key supply chain at stores and warehouses. Among other concerns taken into account were the need for the algorithm to be conceived with practicality for the stakeholders in mind and that it could potentially use sales forecasts as inputs instead of previous year sales. Another important need was for the algorithm's underlying logic to be well documented and adaptable from the outset. If improvements to the process led to more frequent allocation reviews and therefore shorter target periods or if major supply-chain costs like the shipping or picking costs change, the model should be able to accommodate these alterations easily through simple input change.

The third phase was the progressive implementation of what was determined in phase two. Using the process maps defined in the first phase, every cost proxy was progressively tailored in workshops with the different stakeholders. A large portion of time was also dedicated to data extraction and transformation, because the large number of proxies needed for the model required data from multiple information sources spread across different databases of the company.

The path to take the developed theoretical model to a stage where it could be used as a business application required that the prototype developed in the context of the thesis bared some resemblance to the intended final tool to be implemented.

Needed data is stored by the retailer in 2 separate databases and spread across a series of data tables. Before any modeling takes place data extraction, transformation and loading must take place. This refers to large and changing inputs like sales forecasts, PS values, etc. Querying was done through a series of purposefully build SQL queries, that generated .txt files which were then used to communicate with a *Python* script , *cleaning.py*, responsible for table cleaning an aggregation. The script generates a single table that is automatically loaded into *model.xlsl*, a Microsoft Excel ©file, where all the Supply-chain cost modelling takes place. In the same excel file the user is asked to insert some required manual inputs, mostly related with restriction capacity, pallet and container ship utilization rations and the main costs of preforming tasks. The decision to

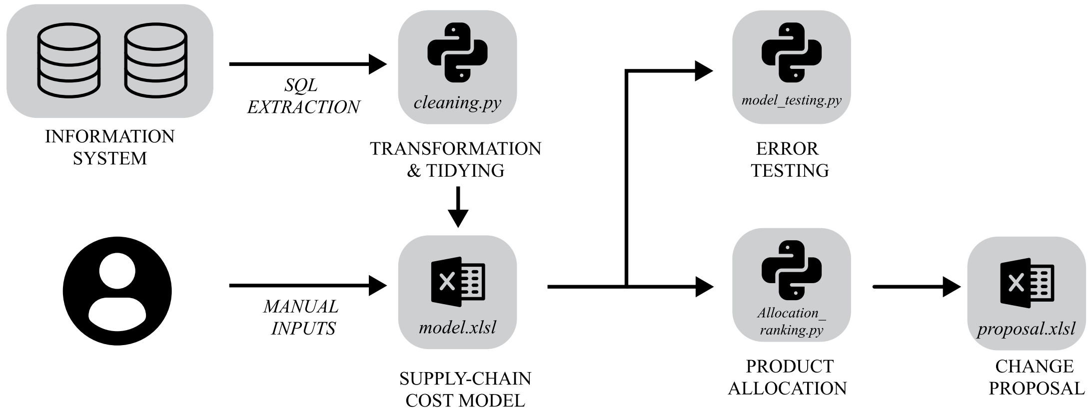

Figure 5.2: Structure of the MVP

use Microsoft Excel for the cost modelling component instead of *Python* was due to the necessity to quickly and easily change proxies as a result of the successive meetings being held and to the need for transparency on the way the cost modelling worked (as most of the involved stakeholders are not familiar with the use of *Python* this could unintentionally become a black-box).

After cost modelling is preformed, a final table with the main outputs of the model is used as the input for the second phase (the allocation heuristic) that runs in a separate python script, called *allocation_ranking.py* The output of this script is again exported as an excel and file (*proposal.xlsl*) containing the final recommendation of which SKUs to reassign.

An additional module is the *model_testing.py* that compares the results of the model with known historical data. It was created to test the feasibility of the model, and only works when the data being inputted to the model is historical data. This component is of crucial importance both to detect errors in the modelling and to convince stakeholders of the validity of the approach. It's also the source for most of the analysis preformed in chapter 6. A scheme of the MVP components and their interrelationships is displayed in Figure 5.2.

5.3 Retailer's Input Data

In this last section of the chapter some notes are made about the data inputted in the model that originated the results discussed in chapter 6.

As the goal of the project was set at defining a minimum viable product (MVP) the choice was made to test the model with historical data, instead of the intended use of sales forecast, which will be their final input once in normal daily use. This allowed a comparison between model predictions for average inventory levels and actual average historical inventory data (at least for the products that historically were assigned to a certain flow). This both provided the needed trust in the model's effectiveness and allowed for manual adjustments to proxies if needed.

The target period chosen for the study was the 2 first months of 2021. Table 5.2 displays the adapted inputs and table 5.3 shows the historical data used as reference.

Table 5.2: Data inputs used as a replacement for the data that will be used in normal operation

Proxy-data, at a product i and store s level		**Input in normal operation , at a product i and store s level**
Average historic daily sales, during the target period	$\rightarrow$	Forecast average daily sales
Average standard deviation of daily forecast errors (in % of forecast sales) 10 day sample of the most recent data	$\rightarrow$	Average daily forecast error (in % of forecast sales)
Current product MDQ	$\rightarrow$	MDQ at the intended period
Current product PS	$\rightarrow$	PS at the intended period

Table 5.3: Algorithm outputs and the historical data that is used as a comparison

Verification data at a product i and store s level		**Output to be verified**
Average historic inventory, at a product i and store s level, during the target period	$\rightarrow$	Predicted inventory, at a product i and store s level, during the target period
Average historical inventory, at a product level i, in Água de Pena warehouse, during the target period	$\rightarrow$	Predicted average inventory at the regional warehouse, at a product level i

Cost drivers and constraints

The two biggest deciding factors between the centralized and non-centralize replenishment are on the one side the logistical costs (in the amount of operations preformed, the cost differentials between warehouses and transport efficiencies) and on the other the differences in replenishment periods, reflecting longer order lead time and review periods. These drive the inventory, zero-inventory stock-outs and in-store shelf replenishment up. All of these points largely align with the findings of Holzapfel et al. (2018).

Logistical cost drivers

For confidentiality reasons exact operation costs aren't disclosed in this thesis, but a comparative analysis can be made: In terms of box handling and picking, the regional warehouse presents better efficiencies than the central one, translating to lower costs per box. Transportation to the regional warehouse is much more efficient in single product pallets, of the centralized flow. Single product pallets can also be stacked on top of each other (something not possible with PBS pallets), translating to an higher multiplier of pallets per shipping container and are consider to have a useful volume ratio of 100%, whereas for PBS pallets a reference value of 80% of volume utilization was considered.

Replenishment Periods

Trucks with shipping containers leave the warehouses the afternoon before the container ship sets sail. There are 2 boats per week leaving the Port of Lisbon to Madeira, with variable trip duration. For the X-docking operation at *AP*, two days are considered as PBL and fresh products are given priority in relation to PBS pallets, because they are for the most part perishable products. In reality some PBS pallets might be expedited within a day.

As can be seen in Figure 5.3 there is a variable replenishment duration in the two weekly delivery schedules. Additionally, stores are allowed to order everyday and not just in the last day before the boat leaves, although these orders are effectively belated until this point.It was decided that integrating these details would make the model more realistic but would result in unnecessary extra complexity for the expected gain in precision to be obtained.

The simplified order lead times and replenishment periods are presented in Figure 5.3.

As consequence of what has been said, the values for R and R_{Wh} were both established to be 4 days and lead-times L and L_{Wh} where established to be 5 and 7 days respectively.

Model constraints

The factor that most restrains the supply-chains capacity is the space at *AP* available to store centralized items. Two types of storage indicators need to be observed. Picking positions, which are ground level positions used by operators to pick products and Reserve positions, where extra stock is held to replace the pallet in the picking position when this one runs out. It should be observed

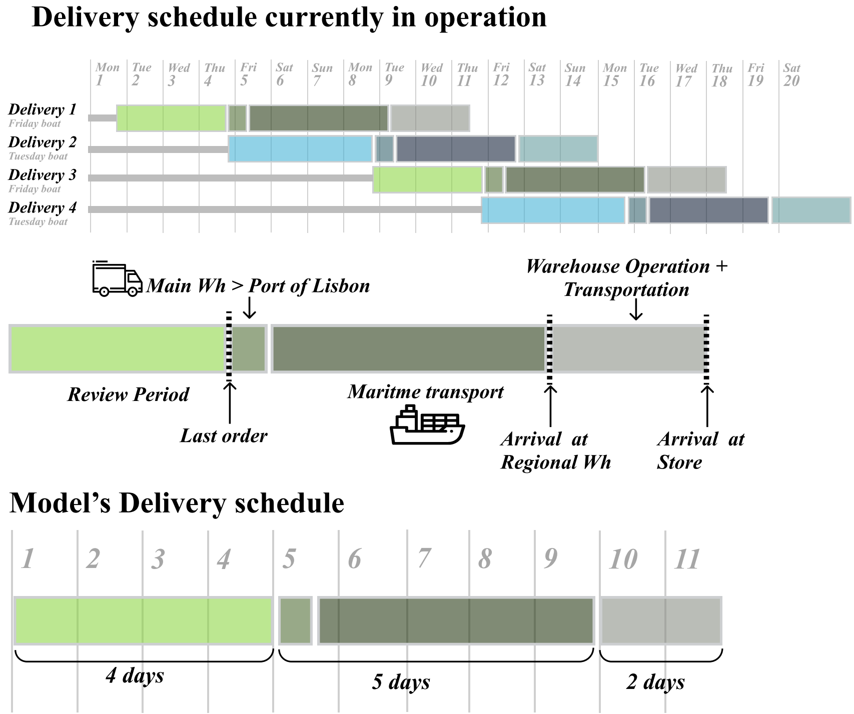

Figure 5.3: Actual and modeled replenishment schedules

that the retailer does not need to use all of its available capacity on the operation, assigning a series of picking and reserve spaces specifically to local suppliers and to extra incoming stock, associated with promotional campaigns. The greedy allocation algorithm is flexible in supporting these changes. For the trial run of the algorithm, a reasonable value close to the value of occupation observed during the target period was used as model inputs. The maximum capacity of picking positions was set at 504 picking positions, and the reserve capacity was set at 2500 reserve positions.

Chapter 6

Results

This chapter's goal is to present the results obtained from the developed tool. It is divided into 3 sections, the first one corresponding to the demonstration for the cost modelling components using examples of specific products (6.1). This is followed by the global allocation recommendation which is the output of the allocation algorithm phase(6.2). Finally, the focus shifts to an analysis of deviations, between model predictions and historical records (6.3).

6.1 Example of cost modelling for a single product

The cost modelling tool crunches the input data for each SKU in the study, predicting costs for each supply chain cost group. Let's dive into the analysis of the alternative flows for a couple of distinct products to get a hint of the major cost drivers. The products chosen as an example were a SKU of bottled soft drink, characterised for being a fast mover, that will be referred as product A, and a SKU of Fruit Juice, which is a slow mover product. These were selected because they present similar physical dimensions but very different demand patterns.

Figure 6.1 presents the costs predicted by the model when supplying products using each of the flows. They also show each component's weight in the total supply chain cost. Note that scales of the axis are not the same for both products.

Naturally, these two products only provide a very narrow understanding of what the global allocation solution encompasses. However, they reveal some interesting insights. As current operations costs stand, transportation costs equate to very large portion of total supply chain costs, due to the required and costly maritime transport. For the two products analysed, they also present the largest possible savings gap achievable by having the SKU allocated to the regional warehouse.

Warehousing costs and extra-handling of boxes also represent some cost savings opportunities, especially in the cost of extra-handling that, for these two products, is even expected to not exist if the centralized flow is used. In practice this means that if these products are centralized, on average, when orders are received there is enough shelf space to accommodate all on-hand inventory. Zero-stock stock-outs exist in both flows and the model indicates that they present lower levels

Product A	***Fast-moving product***					
	Transport	***Warehousing***	***Inventory***	***Stock-outs***	***Extra-handling***	***Total SSC***
Centralized	*4 779,11 €*	*411,31 €*	*22,86 €*	*0,98 €*	*11,80 €*	*5 226,06 €*
Non Centralized	*7 211,82 €*	*468,64 €*	*91,33 €*	*0,49 €*	*529,72 €*	*8 302,00 €*
Change	*2 432,71 €*	*57,33 €*	*68,47 €*	*- 0,49 €*	*517,92 €*	*3 075,95 €*
%	*33,73%*	*12,23%*	*74,97%*	*-100,00%*	*97,77%*	*37,05%*

Relative weights of the cost components in the Total SCC

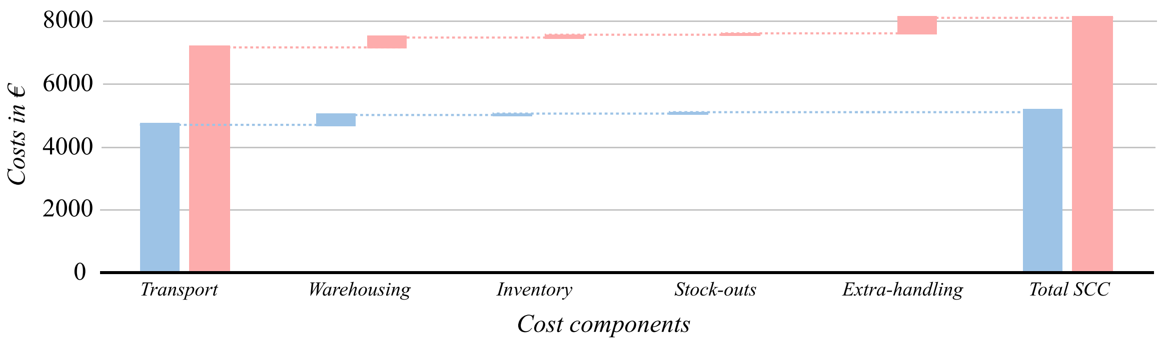

Product B	***Slow-moving product***					
	Transport	***Warehousing***	***Inventory***	***Stock-outs***	***Extra-handling***	***Total SSC***
Centralized	*125,56 €*	*13,03 €*	*14,79 €*	*0,37 €*	*1,18 €*	*154,95 €*
Non Centralized	*178,92 €*	*15,46 €*	*10,09 €*	*0,20 €*	*15,64 €*	*220,30 €*
Change	*53,36 €*	*2,43 €*	*- 4,71 €*	*- 0,17 €*	*14,46 €*	*65,34 €*
%	*29,82%*	*15,69%*	*-46,65%*	*-88,76%*	*92,46%*	*29,66%*

Relative weights of the cost components in the Total SCC

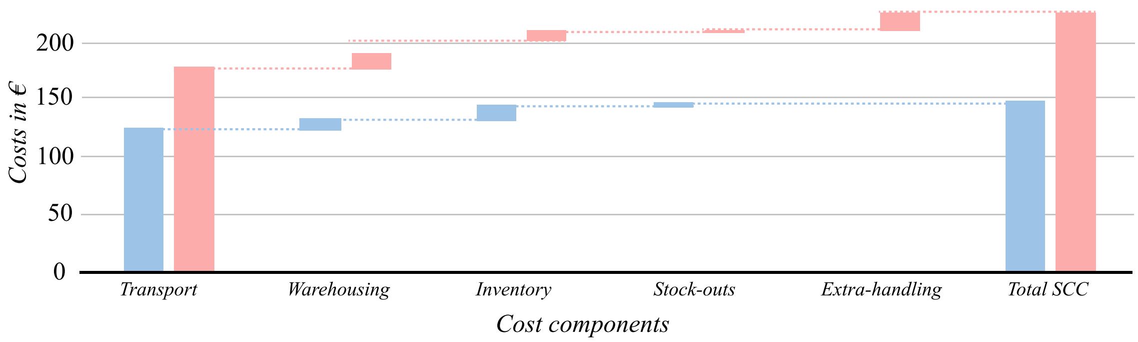

Slow vs. Fast mover contribution

	Slow-mover	***Fast-Mover***
Centralized	*154,95 €*	*5 226,06 €*
Non Centralized	*220,30 €*	*8 302,00 €*
Change	*65,34 €*	*3 075,95 €*
%	*29,66%*	*37,05%*

Non Centralized Cost

Centralized Cost

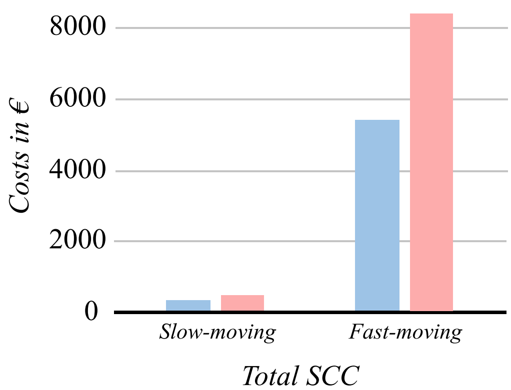

Figure 6.1: Supply Chain Costs of for each flow - Slow-moving and Fast-moving examples

for products that are not-centralized. The difference though is negligible and doesn't represent the full picture regarding stock-outs, as is discussed in Chapter 7.

Inventory costs favour different allocations depending on the product. For fast movers, the pooling of demand in the regional warehouse reduces the overall stock in the region. For slow movers, the additional inventory in the regional warehouse actually offsets the gains of inventory pooling.

It can be concluded that the majority of SKUs benefit from being allocated to the regional warehouse. Observations made above are generally extendable to the whole range of analysed products, with some variation regarding the relative proportion of each of the cost components depending on the physical characteristics of the SKU. If warehousing constraints in the regional warehouse are not considered, the decision of holding stock of an item in the regional warehouse comes down to the balance of each cost component. For fast and medium-fast moving products, savings with zero-stock stock-outs and inventory cost are negligible when one considers the total supply chain costs.

The less the products sells, the less expressive warehousing and logistical costs and savings opportunities are, making inventory and stock-out costs slightly more relevant. As such, for the slow and very slow moving products, the model sometimes indicates that products should be centralized and sometimes non-centralized. However, due to the very slim difference between Total Supply Chain Cost for both flows and the model's level of imprecision (that will be addressed next), the recommendation for these items should be interpreted as indifference on whether to be assigned to a flow or another.

Another important insight is the difference in savings potential between items, as presented in figure 6.1. Although both products would be supplied more cost effectively if held at the regional warehouse, the savings potential of the fast moving product chosen is 47 times greater than the ones of the slow moving product. This shows why the correct prioritization of which products to centralize is so important, given the limitations in the regional warehouse.

6.2 Centralization Recommendation and Benefit Capture

This section addresses the second phase of the allocation approach, that is the result of the greedy algorithm that actually proposes the final assignment of products accounting for warehouse constraints. The first part of the section addresses the proposed changes in flow compared to the real historical allocation. The second part focuses on the effect of the incremental assignment of SKUs to reach deeper insights on the strategic importance of good product allocation.

Current Allocation vs. Proposed Allocation

Figure 6.2 presents the several supply chain costs for the allocation of products during the period studied and what the costs would have been if the allocation suggested by the model was used. In a normal situation, this table would effectively present the potential realisable savings of each cycle of allocation change.

Note that when the tool is in normal functioning, it will be fed by sales forecast data and not historical sales data. To evaluate the accuracy of the tool in a meaningful way the costs of the current allocation are the result of running the cost function with the current allocation and not the actual historical costs incurred in the period. The inventory level predictions for stores and for the regional warehouse are also based on the model's calculations and not on the actual historical data of inventory levels.

This means that biases resulting from the model's proxies will be reflected in both the current allocation cost and the proposed allocation cost, making the comparison between the two proposals fair. After, the model's prediction errors are analysed to identify these deviations in the proxies.

The values presented next should be taken with caution because of (i) the model's deviations, (ii) because some of the most influential cost inputs (such as the number of pallets per shipping container) used are rough approximations, as definitive values were not reached within the duration of the project and (iii) because some of the mandatory allocations to the regional warehouse that are consequence of fiscal and strategic reasons are ignored. One of the estimated costs that deserves specific attention is the inventory cost estimation of centralized products, in the current allocation, as costs are grossly over-estimated as indicated in figure 6.2. [1] Bearing this in mind, promising potential savings of 40563,70€ could have been captured if a better allocation had been used just during the two month study period (02/01/2021 to 01/03/2021). The main drivers of these savings clearly being a more efficient transportation, although all supply chain costs drop, with the exception of the zero-stock stock outs that register a slight increase.

It might seem counter-intuitive to realize that globally costs associated with centralized products actually increase in the proposed solution, which could be seen to be at odds with what was previously stated. This happens because although when looking at a product individually the centralized flow is generally more cost effective than the non-centralized flow, products assign to the centralized flow tend to represent a bigger volume of sales. As more sales volume can be directed through the centralized flow, the bigger this component gets the more savings potential is captured.

Parallel to the cost savings, other improvements in operational indicators should be mentioned, even though they are already partly captured in the cost proxies used. The first is the usage of reserve positions in the regional warehouse, that can now provide a better service to the supply chain stakeholders running at a lower average occupation, through the release of 227 reserve positions. The released rack space represents extra slack that can be used to better cope with peaks of demand of promotional seasons, where extra capacity might be needed. The second is the release of space from the back-room of the stores. The data shows, for a single store that was considered representative, a reduction of 881 boxes to the average number of boxes on-hand in the backroom,

[1] The inventory cost of current allocation was discovered to be largely over-estimated due to the over-estimation of inventory in the regional warehouse. Centralization of products should diminish the total inventory in the region, due to the pooling of demand. However, if the product is a slow mover, a minimum inventory of one full pallet in the regional warehouse might actually make the total region inventory in the centralized flow higher than the total region inventory in the non-centralized flow. This artificially increases the inventory costs. Some of the items currently held on the regional warehouse are slow moving, high value density goods such as spirits and other alcoholic beverages. Slight over-estimation of inventory in these items cause big increases in the inventory. Because of this current inventory costs and inventory cost savings described are not reliably modeled.

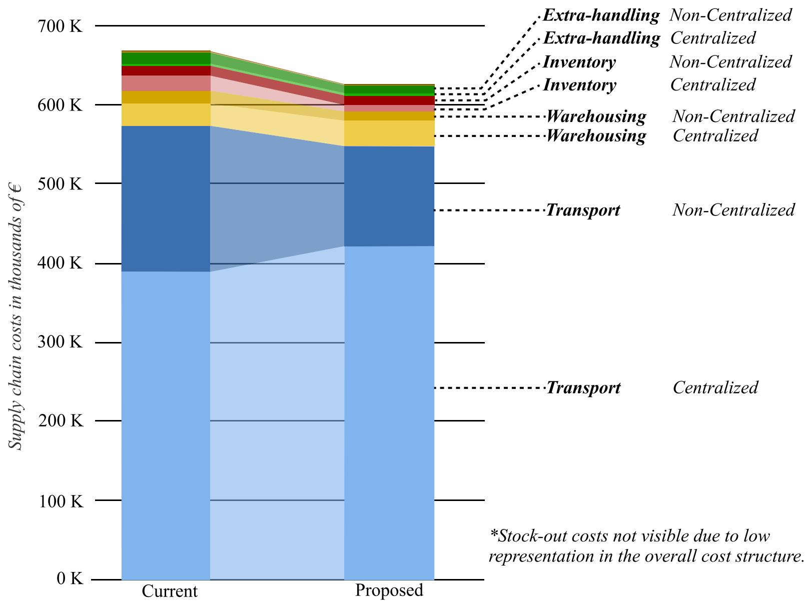

Other Operational Indicators

		Picking positiosns	Reserve positions	Store backroom boxes
Current allocation	Centralized	504	1495	194
	Non Centralized			3675
	Totals			3869
	Utilization		44,96%	
Proposed allocation	Centralized	504	1268	234
	Non Centralized			2754
	Totals			2988
	Utilization		38,14%	

Current Allocation Costs vs. Proposed Allocation Costs

		Nr of SKUs	Transport	Warehousing	Inventory	Stock-outs	Extra-handling	Total SCC
Current allocation	Centralized	504	390 439,66 €	28 795,36 €	* 18 547,83 €	1072,396001	3 524,66 €	442 379,91 €
	Non Centralized	2089	184 845,42 €	15 702,88 €	11 987,22 €	455,153986	13 765,92 €	226 756,60 €
	Totals		575 285,08 €	44 498,25 €	30 535,05 €	1527,549987	17 290,58 €	669 136,51 €
Proposed allocation	Centralized	504	422 944,59 €	31 363,98 €	8 953,57 €	1368,263063	4 144,16 €	468 774,55 €
	Non Centralized	2089	126 886,25 €	11 733,82 €	10 881,60 €	311,7624398	9 984,78 €	159 798,22 €
	Totals		549 830,84 €	43 097,80 €	19 835,17 €	1680,025503	14 128,94 €	628 572,77 €
		Nominal value	- 25 454,24 €	- 1 400,45 €	- * 10 699,88 €	152,48 €	- 3 161,64 €	- 40 563,74 €
		% Current vs. Proposed	-4,42%	-3,15%	-35,04%	9,98%	-18,29%	-6%
		% Proportion of Savings	-62,75%	-3,45%	-26,38%	0,38%	-7,79%	* See footnote [1]

Figure 6.2: Current allocation vs. Proposed allocation

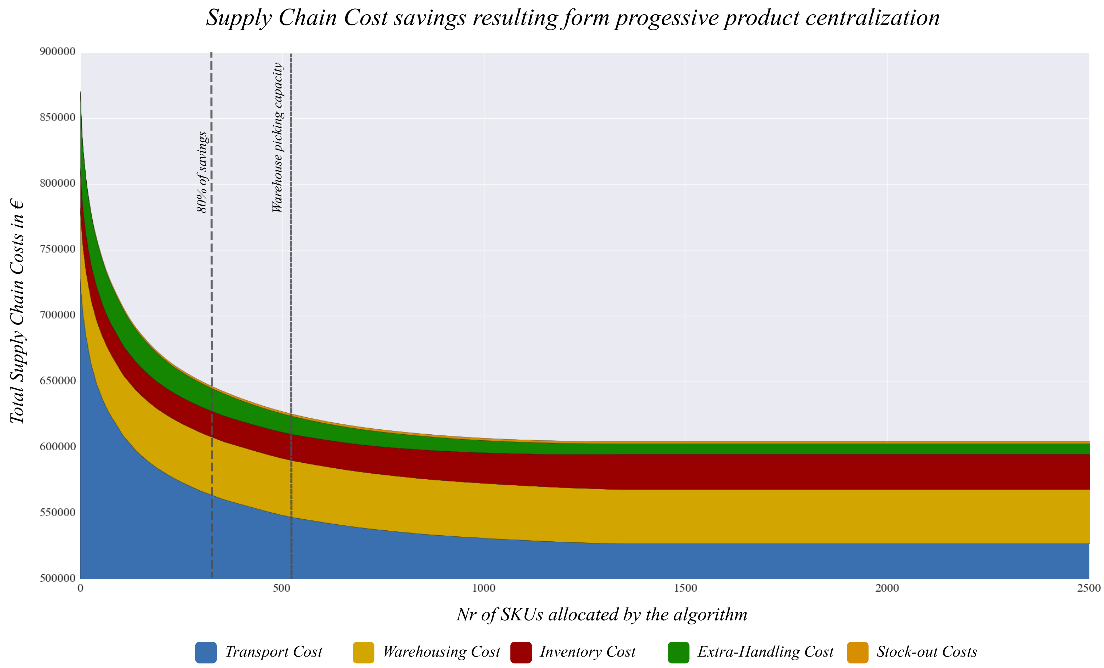

Figure 6.3: Savings potential Curve

representing a 23% decline.[2] This is very important given extra pressure put on stores backrooms in the region, due to the extra long replenishment periods.

Other Allocation Related Insights

The chosen approach used for assigning the SKUs to the regional warehouse allows for some interesting further study of the marginal impact that progressively assigning items to the regional warehouse has. As mentioned before, products are assigned in order of their individual savings if they are centralized. Figure 6.3 was generated by tracking the predicted total supply chain costs as the algorithm progressively assigns more products to the regional warehouse. Warehousing constraints were removed to understand what would be the movement of the costs savings even after the limit imposed by the warehouse would be reached. Notice that while previously "savings" was interpreted as the comparison between total supply chain cost difference between the current allocation and the proposed one (translating the realisable savings), in this subsection savings have another interpretation: they translate the savings in total supply chain cost between the proposed allocation at any given step the algorithm is calculating and the option of having all products in the non-centralized flow (as this is the baseline where the algorithm calculation starts).

Savings are heavily concentrated in the first products that are assigned. Of the possible 265702,88€ of savings resulting from centralization during the target period (if the warehouse was unconstrained), 80% can be captured with the assignment of 311 products or 11,98% of all

[2]The decline refers only to boxes associated with this flow and not to an drop in the total number of boxes

SKUs include in the scope. This value is well within the constraint of the picking positions. Both the picking positions capacity and the mark of 80% of savings are displayed with a dotted line in Figure 6.3.

Additionally, and in line with the above observations made for the single product cost model, the biggest driver of savings is by far the transportation component.

6.3 Analysis of model errors

The cost model calculates supply chain costs based on the assumptions and simplifications of the inventory policy used. As with any simplification there's a risk of imprecise predictions that then lead to incorrect cost modelling and allocation. The model can face two types of deviations. Deviations caused by the proxies used will produce errors in predictions for all products and locations. Deviations associated with individual products, either by errors in the parameters stored in the retailer's information system or by manual overrides of orders, done by the companies commercial managers, produce errors only on these specific products.

The first step to be taken regarding this problem is the quantification of the errors produced by the model and by understanding if the retailer is comfortable with the achieved level of precision. This relates also with the relative weight of each cost proxy in the final cost function and the impact the errors have on them.

After these assessments corrective actions can be taken. For the first type of deviation, possible actions include improving the cost proxies or manual offsets that artificially force results to be closer to historical records. For the second type of deviations, outlier values should be cleaned and the root cause for the excessive deviation should be investigated (e.g. by identifying common characteristics in these products).

The analysis of errors for products of a flow can naturally only be done for items that historically have used that flow. Therefore the analyses always appear divided between products assigned to the regional and the central warehouses.

The analysis of deviations focuses on two key metrics generated by the model: average inventory at stores and average inventory at the regional warehouse . These were chosen because of the availability of historical data and because most proxies are in practice multiplied by these values. (i.e. if these are proven to be accurate the model as a whole should be accurate too, assuming cost inputs are also accurate). The analysis can be seen in Figure 6.4.

Store Inventory Comparison

The store inventory comparison is presented both in percentage deviation between predicted and historical inventory and in absolute number of boxes, for the whole region. This is done because the proportion metric in products that sell very small quantities is much more sensitive, while not representing a big impact in store operations. Note that deviations are presented as the sum of all stores in the region which is an abstraction that doesn't transmit to well the situation at each store.

	Whole Region				Single store	
	Centralized		Non-Centralized		Centralized	Non Centralized
NR of SKUs	504		2089		**500**	**1953**
	(1) **Deviation (%)**	(2) **Nr of boxes**	(1) **Deviation (%)**	(2) **Nr of boxes**	**Nr of boxes**	**Nr of boxes**
Mean	-63,53%	-112,56	25,71%	0,95	-73,87	-2,08
Std	20,27%	240,24	299,28%	35,53	129,71	91,25
Min	-98,18%	-2782,20	-98,81%	-487,90	-1149,00	-3554,00
1st Quartil	-77,59%	-104,32	-42,76%	-8,64	-85,00	-12,00
Median	-66,52%	-48,84	-13,06%	-1,94	-31,00	-2,00
2nd Quartil	-53,76%	-21,12	41,51%	7,90	-11,00	9,00
Max	35,75%	4,42	7900,00%	758,03	32,00	424,00

(1) *for each product i* $\left(\sum_{s=1}^{RegStores} E(OHinv_{i,s,f}) - \sum_{s=1}^{RegStores} HistoricOHinv_{i,s,f}\right) / \sum_{s=1}^{RegStores} HistoricOHinv_{i,s,f}$ (2) *for each product i* $\sum_{s=1}^{RegStores} E(OHinv_{i,s,f}) - \sum_{s=1}^{RegStores} HistoricOHinv_{i,s,f}$

Figure 6.4: Store inventory error analysis

Because of this, box deviation data for a single store considered representative is also displayed in the results table.

The main conclusion taken from the deviation analysis is that the cost model is prone to deviation when it comes to predicting On-Hand inventory in stores. Historically centralized products had much less on-hand inventory than the model predicted. Non-centralized products had predictions with deviations but these were centered around 0, with over predictions for some products and under predictions for others.

Several reasons can be found to explain this disparity. Not all products follow the replenishment policy explained in the course of this document. A great portion of the items follow a min-max replenishment policy set out by individual product group managers, belonging to the commercial department. This however called for a separate analysis for products that follow this system and that will be from now on referred to as ENA products, as this is the information system program used to manage their orders. Products replenished using this policy are distinguished in Figure 6.4 by a different color. The decision to model the replenishment policy by approximating it to an (R,S) standard policy even though many products don't use this system was twofold. Firstly, the retailer intends to switch most of the products sold in the region to this method of replenishment in the near future. Secondly, this method provided a relatively reliable and standard benchmark for all products that could be implemented into the tool more easily than the many different replenishment parameters set for products that don't use this replenishment method.

Other relevant reasons for prediction deviations can be manual orders placed by commercial managers, the simplifying assumptions of the cost model (e.g. constant demand, no cap on orders, no consideration for store-pack size) and errors in the input parameters of some of the SKUs. This is mostly related with poorly define values for Presentation stock, Minimum Display Quantity and box volume.

Regional Warehouse Inventory Comparison

Deviations for regional warehouse stocks are presented in absolute number of pallets per product i. This point is particularly important because this is the main bottleneck of the whole system and the limiting constraint in the allocation algorithm.

Here the cost modelling preforms relatively poorly especially for the products that sell the most, predicting much less pallets than the actual historical quantity held at the warehouse. For 14 products there where more than a 20 pallet difference between model predictions and historical data. These deviations might put in question the proposed allocation, as there is the possibility the selected SKUs surpass the available capacity in the regional warehouse. As this is such a critical point, it might justify the redesign of the warehouse inventory proxy in future work. Figure 6.5 displays the data related with deviations between prediction of warehouse positions used and historic usage of warehouse positions.

Nr of Products	*501*
Mean	*0,37*
Min	*7,9*
25%	*-4*
50%	*-1*
75%	*1*
Max	*4*
25%	*64*

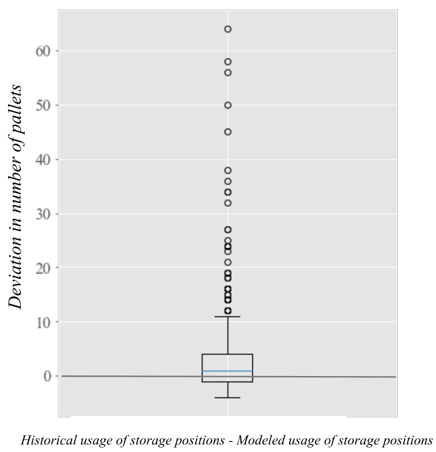

Figure 6.5: Deviations in warehouse rack utilization

Chapter 7

Conclusion and Future Work

The last chapter closes the document by commenting on the achieved results (7.1), discusses some of the cost elements left out of the approach (7.2) and ends with an overall reflection on the project and suggestions for future work (7.3)

7.1 Results Analysis

The achieved results, although quite promising need to be contextualized, due to the deviations found between model prediction and historical records. Use of the model to extract an actionable proposal for altering product allocations should only be done after further improvements and testing against historical data. Even with these shortfalls, two very important points need to be made.

The first is the overall robustness of the underlying logic used in the cost modelling and the allocation prioritization. A model was created that captures the overall logistic operation to a good degree and that was the biggest goal set out for the project. The fact that mitigation of model errors can be done by mostly focusing on adjustments to the defined cost proxies and not through major changes to the way they are calculated (except possibly to the proxy for warehouse inventory) is a proof of that.

The other major point is that the direction of errors and the tendency for diminishing returns when progressively assigning products means that although savings estimations might be off, the algorithm is still proposing an allocation mix that is very efficient. For once, transportation costs, that are by far the biggest cost component, are unaffected by deviations in inventory predictions at the store level. On the other side, the model's tendency to over-estimate inventory for centralized flow and under-estimate it for the non-centralized flow, means that cost savings stemming from on-hand inventory at stores and extra-Handling would actually be even greater. Finally, the sharp diminishing returns in savings of progressively centralizing the next product mean that, even if the centralizing ranking that the algorithm uses would change slightly with more accurate proxies, most of the savings would be captured anyway, because "top-savers" would be featured in the first places of the ranking either way.

The model takes some simplifications that do deviate slightly from the actual operation. One such points is the need for some products to be held at the regional warehouse regardless of the supply-chain cost in either flow. These default allocations can have their origin in the retailer's own strategy or on external regulatory issues. The model can be easily altered to account for these products. It can also be useful in quantifying the additional cost of using these default allocations.

7.2 Out of Scope Factors

In the first phase of the approach only costs that were directly influenced by the allocation and whose influence could be reliably translated by a cost proxy were included in the model. This left out potentially important cost factors such as the previously referred stock-outs with stock (i.e. when there is stock in the backroom of the store that for some reason fails to be replenished in the shelf). While the project was underway, it was asserted that non-centralized products experienced higher levels of stock-outs with stock than centralized products, while stock-outs without stock were fairly similar for both flows (as predicted by the model). If included in the model this could potentially affect the Total Supply Chain Cost predictions has it places greater importance on stock-out costs. However, because it was not possible to determine the root cause of the stock-outs, linking them directly with the allocation decision, a cost proxy for stock-outs with stock was not considered for the work here presented. Integration of this cost in the cost modelling is one of the main tasks to be included in the project's suggested future work.

7.3 Concluding Reflections and Future work

The developed work was successful in proposing a new holistic view to revamp a legacy process, showing potential for both savings and operational improvements. It clearly demonstrates a situation where optimization had to take the complete internal supply chain into account and not rely only on optimization of individual supply chain subsystems.

The projects most relevant deliverable took the form of an optimization strategy consisting of a cost modelling components and a prioritizing heuristic. This strategy is flexible enough to account for possible future changes in the supply chain's cost structure. This strategy materialized itself in the form of a prototype handed to the retailer to serve as the basis for the development of the final form of the tool.

Together with the strategy and the prototype, another valuable output of the project were the important insights obtained. The knowledge of the relative weights of the individual cost factors, namely of the transport cost, is useful in directing the focus of future improvement efforts. The project revealed that the retailer could have captured substantial savings during the studied period, if a different allocation had been used.

The hope is that this new tool, if adopted by the stakeholders and used with due frequency will allow for the capturing of savings associated with correcting these inefficiencies. Future work should focus on two axis. The first is the improvement of the model's accuracy. This can happen by refining the proxies used, namely the predictions of on-hand inventory in stores and on the regional warehouse, by cleaning errors in the input data and by manually tweaking some proxies to better match historical records. The second axis is the focus on the allocation tool's agility, practicality and ease of use. Integrating the tool with the retailer's information system, and packaging the several developed components into one cohesive tool are some of the tasks that form this second axis. Future work might also include the expansion of the analysis to the frozen products warehouse, That function in parallel with the ambient temperature studied.

Bibliography

de Koster, R., Le-Duc, T., and Roodbergen, K. J. (2007). Design and control of warehouse order picking: A literature review. *European Journal of Operational Research*, 182(2):481–501.

Eroglu, C., Williams, B. D., and Waller, M. A. (2013). The backroom effect in retail operations. *Production and Operations Management*, 22(4):915–923.

Fernie, J. and Sparks, L. (2018). *Logistics and Retail Management: Emerging Issues and New Challenges in the Retail Supply Chain*. Kogan Page.

Fréville, A. (2004). The multidimensional 0–1 knapsack problem: An overview. *European Journal of Operational Research*, 155(1):1–21.

Holzapfel, A., Kuhn, H., and Sternbeck, M. G. (2018). Product allocation to different types of distribution center in retail logistics networks. *European Journal of Operational Research*, 264(3):948–966.

Hübner, A. H., Kuhn, H., and Sternbeck, M. G. (2013). Demand and supply chain planning in grocery retail: an operations planning framework. *International Journal of Retail & Distribution Management*.

Jacobs, F. and Chase, R. (2012). *Operations and Supply Chain Management: The Core: Third Edition*. McGraw-Hill Higher Education.

King, P. (2011). APICS magazine (August) 'crack the code'.

Lapinskaitė (2014). The impact of supply chain cost on the price of the final product. *Business, Management and Education*, 12:109–126.

Lovell, A., Saw, R., and Stimson, J. (2005). Product value-density: managing diversity through supply chain segmentation. *The International Journal of Logistics Management*.

Martello, S. and Toth, P. (1990). *Knapsack Problems: Algorithms and Computer Implementations*. Wiley Series in Discrete Mathematics and Optimization. Wiley.

Rushton, A., Croucher, P., and Baker, P. (2014). *The Handbook of Logistics and Distribution Management: Understanding the Supply Chain*. Kogan Page.

Silver, E., Pyke, D., and Thomas, D. (2016). *Inventory and Production Management in Supply Chains*. Taylor & Francis.

SONAE (2021a). Apresentação da sonae mc.

SONAE (2021b). História do grupo.

SONAE (2021c). O grupo e os negócios.

Sternbeck, M. G. and Kuhn, H. (2014). An integrative approach to determine store delivery patterns in grocery retailing. *Transportation Research Part E: Logistics and Transportation Review*, 70:205–224.

Tompkins, J., White, J., Bozer, Y., and Tanchoco, J. (2010). *Facilities Planning*. Wiley.