A Simulation-Optimization Model to Determine Fashion Delivery Patterns

Daniela Timóteo Martins de Pinho

Master's Dissertation

Supervisor: Prof. José Fernando da Costa Oliveira

U. PORTO

FEUP FACULDADE DE ENGENHARIA UNIVERSIDADE DO PORTO

Mestrado Integrado em Engenharia e Gestão Industrial

2019-07-01

Abstract

Fashion retailers must decide on the best strategy to deliver products from warehouses to stores, maximizing the fulfillment of demand and keeping costs low, while simultaneously dealing with short life cycle products, a high risk of lost sales caused by stock-outs and backrooms with limited capacity. This industry's specific challenges have not yet been fully researched regarding the definition of an ideal delivery pattern, which is capable of quickly replenishing sales and deliver new products to the stores, filling up shelves after the weekend demand peak.

The simulation-optimization method proposed leans on simulation to generate sales and shipping data, providing inputs for optimization which are not tied to past conditions. This requires a prior step where demand is estimated for days where its complete fulfillment was not possible, and, therefore, not recorded, due to stock-outs. Owing to the inherent complexity of fathoming the impact on sales of different delivery days, and in order to avoid the computationally intensive simulation of the entire solution search field, a metamodel is introduced to predict the impact on store profits of different replenishment speeds. Afterwards, the linear programming problem is solved generating the delivery pattern which better balances service level and logistic costs. The holistic integration needed for processing in the distribution center is assured by including constraints tackling operational limitations. The best alternative is then tested at the SKU-store-day level on the simulator and a detailed solution is obtained for each store, allowing validation.

By introducing a case study on an Iberian fashion retailer, the methodology is implemented and results show a 3,0% decrease in overall costs, which incorporates a contributing increase in profits when comparing to the currently implemented policy. Additionally, simulator features allowed experimentation with emergency shipments for stocked-out products, which resulted in an expected 6,9% reduction in costs.

Resumo

Empresas que operam na indústria da moda encaram a definição da melhor estratégia para entregar produtos às lojas, através de armazéns centralizados, maximizando a concretização da procura e mantendo custos controlados. Isto é amplificado pelos desafios que enfrentam como o facto de se tratarem de produtos com um reduzido ciclo de vida, para os quais há um grande impacto em vendas quando surgem ruturas, e ainda limitados pelo espaço de armazenamento em loja. As características específicas desta involvência ainda não foram investigadas em profundidade na definição de um modelo de abastecimento capaz de restabelecer rapidamente vendas e entregar novas coleções nas lojas, enchendo prateleiras após o pico da procura do fim de semana.

O método de simulação-otimização proposto utiliza a simulação para gerar dados de vendas e de expedição, fornecendo *inputs* para otimização que não estão presos a condições passadas. Isto requer uma etapa anterior em que a procura é estimada para os dias em que seu cumprimento total não foi possível, não sendo, portanto, registada, devido à inexistência do produto na loja. Devido à complexidade inerente de avaliar o impacto nas vendas de diferentes dias de entrega, e para evitar a simulação computacionalmente intensiva de todo o espaço de pesquisa da solução, um metamodelo é introduzido para prever o impacto nos lucros da loja de diferentes velocidades de reabastecimento. Posteriormente, o problema de programação linear obtem as combinações de dias de entrega que melhor equilibram nível de serviço e custos logísticos. A integração holística necessária para o processamento no armazém é assegurada pela inclusão de restrições que lidam com as limitações operacionais. Posteriormente, a melhor alternativa é testada ao nível do SKU-loja-dia no simulador e uma solução detalhada é obtida para cada loja, permitindo validação.

Através da introdução de um caso de estudo num retalhista ibérico de moda, a metodologia é implementada e os resultados mostram uma redução de 3,0% nos custos globais, o que incorpora uma diminuição da margem perdida, quando comparado com a política implementada atualmente. Além disso, o simulador desenvolvido permitiu o teste da implementação de um fluxo de urgência para produtos sem stock em loja, o que resultou numa redução esperada de 6,9% nos custos.

Acknowledgements

I remember being on the second year of the degree warning freshmen on just how quickly time flies by, but I guess I never truly understood that five years do, in fact, last only slightly more than a heartbeat. That feeling comes, however, with the realization that I wouldn't have made it as far without the endless support of the innumerous people that have crossed my path.

Firstly, I would like to thank my supervisor Professor José Fernando Oliveira, for his availability and knowledgeable insights. Improvement opportunities were not left unnoticed.

I am grateful for Daniel, Raquel and Diogo, my project team, for all the brainstormings, help and advice you provided, building the final result accomplished. Your work set the baseline for the quality I will always keep striving for.

To Paulo Sousa, who guided me whenever I was left unsure and in need of organizing my strategy. Your careful review of my work heightened my confidence in it.

I am extremely thankful for every LTPeer that wasted no time on making me belong and whose minds never cease to amaze me. You redefined my meaning of colleagues, by creating a work-and-fun environment I could have never predicted.

I must not forget the continuous safe haven my family guaranteed, as having a place to get back to is too often taken as granted. To my dad, who through jokes actually managed to push me to pursue Engineering and who is a true example of perseverance.

And last (but definitely not least), the amazing people I get to call friends. My everlasting 'complicated' girls turned women who possess the special ability of making me hit the reset button on time and remember how to live life as I did on easier times. And the classmates turned much more who never doubted me, even in the too often times I did. Thank you all, for making life worth living and truly inspiring me to keep shaping the best version of myself.

Daniela Pinho

"O futuro?... A construção no presente!"

Contents

1 Introduction **1**
1.1 Motivation . . . 1
1.2 Project description . . . 2
1.3 Dissertation structure . . . 3

2 Theoretical background **5**
2.1 Fashion Industry . . . 5
2.1.1 Short life-cycles . . . 5
2.1.2 High Volatility . . . 6
2.1.3 Low Predictability . . . 6
2.1.4 High impulse purchase . . . 6
2.2 Supply Chain . . . 7
2.2.1 Procurement . . . 7
2.2.2 Distribution . . . 8
2.2.3 Replenishment methods . . . 11
2.3 Demand estimation . . . 11
2.4 Simulation-Optimization . . . 12
2.4.1 Optimization . . . 12
2.4.2 Simulation . . . 13
2.4.3 Metamodel-Based Simulation Optimization . . . 13

3 Problem description **15**
3.1 Case study context . . . 15
3.1.1 Store . . . 16
3.1.2 Transportation . . . 17
3.1.3 Distribution center . . . 18
3.1.4 Costs overview . . . 18
3.1.5 Seasonality . . . 18
3.2 The challenge . . . 19
3.2.1 Current approach . . . 20

4 Methodology **21**
4.1 Proposed approach . . . 21
4.2 Demand estimation . . . 22
4.3 Simulation . . . 24
4.3.1 Model . . . 25
4.3.2 Validation . . . 30
4.4 Metamodel . . . 32
4.4.1 Model . . . 32

4.4.2 Validation 33
4.5 Optimization 34

5 Case Study Results 39
5.1 Delivery Pattern optimization 39
5.2 Emergency shipments simulation 41
5.2.1 Contextualization 42
5.2.2 Results 42

6 Conclusions and Future Research 45

Acronyms and Symbols

3PL	Third-Party Logistics
DC	Distribution Center
MAPE	Mean Absolute Percentage Error
MPE	Mean Percentage Error
OUTL	Order Up-To Level
QR	Quick Response
SKU	Stock Keeping Unit

List of Figures

1.1 Project's timeline. 2

2.1 Fashion management approaches. Source: Martino (2015). 8
2.2 Retail logistics subsystems. Adapted from Kuhn and Sternbeck (2013). 8
2.3 Breakdown of store operations. Source: Pires et al. (2017). 9

3.1 Sales quantities and profits per week for the year 2018. 19

4.1 Main steps of the methodology presented. 21
4.2 Main blocks of the simulation model. 25
4.3 Comparison of sales quantities per week between simulation results and reality. . 31
4.4 Comparison of number of pallets shipped per week between simulation results and reality. 31
4.5 Example of the relationship between replenishment speed and obtained profit for one store. 33

5.1 Comparison of the costs obtained for multiple scenarios tested. 39
5.2 Comparison of the percentage of units over capacity at store reception and DC preparation for multiple scenarios tested. 41
5.3 Comparison of the costs obtained with or without emergency shipments. 43
5.4 Comparison of the average quantities received for each day of the week with or without emergency shipments. 43
5.5 Sensitivity analysis of the impact of changing sorter operating time for the emergency shipments scenario. 44

List of Tables

3.1 Weekly seasonality of global chain sales for 2018. 17
3.2 Costs per subsystem 18

4.1 Table of notation for demand estimation. 23
4.2 Conditions for each type of shipment request. 29
4.3 Table of notation for optimization. 37

5.1 Comparison of the distribution of delivery frequencies for the As-Is and the Optimized scenario. 40

Chapter 1

Introduction

1.1 Motivation

The importance of the fashion industry in the global economy is indisputable, with growth happening every year (Amed et al., 2018) and players becoming increasingly competitive. The large dimension comes with extra layers of complexity, as the market is characterized by short product life-cycles, high volatility, low predictability and a high amount of impulse buying, creating the need for strong availability (Christopher et al., 2004). The capacity to present the right product, at the right time, at the right place represents a greater fulfillment of demand, in line with less stock-outs and, consequently, higher sales and profits, with fewer markdowns happening at the end of the season.

In the last years, these challenges have been reflecting on a change from a pure push model to one pulled by demand, as fast-fashion concepts arose leading to the definition of multiple periods within a season and the constant introduction of new products in stores. This creates the need for a paradigm shift in the way stores are supplied, as permanent collections become scarcer and the need for replenishment of sales needs to be balanced by the delivery of new products. Postponement, by the use of a central warehouse to delay the distribution of stock, allows for a shorter time-frame for forecasted demand at the store dimension, mitigating errors on the predictions made at the SKU level when the product orders were made to the supplier. Having a warehouse also goes in hand with the limited backroom capacity some stores possess. Understanding the ideal store delivery frequency can provide cost reduction possibilities by finding the accurate trade-off between demand fulfillment and associated costs.

The latest report presented by McKinsey&Company (Amed et al., 2018) states that as economic conditions are predicted to worsen in 2019, companies are more than ever focused on reducing costs and increasing efficiency. To do so in a longer time period, analytic-based decision-making is recommended for optimized results. The present project builds on this call, attempting the redefinition of replenishment strategies, by leveraging data-backed models to support strategic business decisions. The challenge of analyzing the complete impacts of different measures can then be undertaken without a trial-and-error methodology.

Motivated by a case study of an Iberian fashion retailer, the current delivery patterns fixed by long-term empirical tweaking can be questioned by the use of analytic tools, enabling the holistic assessment of supply chain implications and demand satisfaction.

1.2 Project description

The stakeholders' awareness of the potential implications of redesigning replenishment practices led to the present study ventured as a consulting project. The company at stake has been consistently expanding by opening new shops and entering new markets to spread its successful business. Since tighter restrictions are posed by shipment possibilities for farther away locations, the scope was limited to centrally-controlled stores within Portugal and Spain, accounting for more than a hundred different facilities.

The project was overseen by a team of entities responsible for all the major dynamics involved, from warehouse processes and day-to-day operations to overall strategy definition. This tight relationship enabled smooth access to information and knowledge gathering, as well as the understanding of the expected outcomes and validity assessment, within predefined business constraints.

As presented in Figure 1.1, the first stage was targeted at collecting all the relevant data, establishing the As-Is process and the degrees of freedom provided. With a strong baseline formed, a strategy was formulated to tackle the problem, by analytically modelling demand and beginning the definition of the simulation model. The overlap between stages allowed the clarification of eventual lacks and undetermined restrictions. The four main elements of the approach were tackled as needs arose, as firstly, demand was estimated to provide clear inputs for the simulation, which then proved incapable of microscopically attempting every combination of delivery days for all the stores, issuing the need for a metamodel and an optimization module. Throughout the formulation of these models, validation was performed, ensuring adherence to reality and reliable results. The last step gathered conclusions and elaborated on the recommendations proposed for implementation.

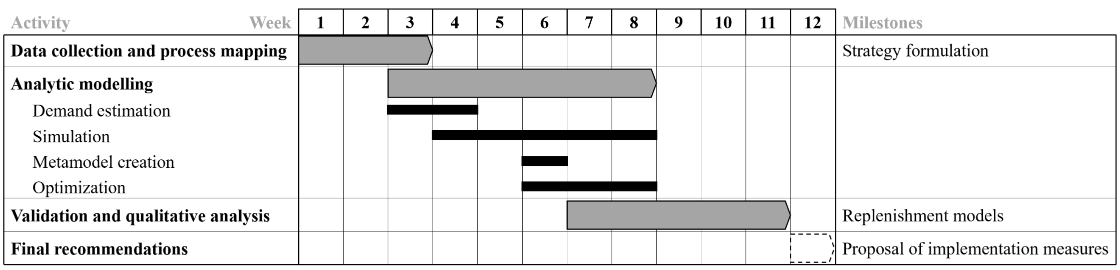

Figure 1.1: Project's timeline.

1.3 Dissertation structure

The organization of this dissertation aims at providing a logical sequence of topics, guiding the reader throughout the formulation of the proposed approach. Chapter 2 elaborates on the current literature of the matter at hand, providing an in-depth contextualization of the problem. Relevant contributions are summarized, establishing the potential for knowledge extension. The third chapter underlines the characteristics of the case study, creating an improvement opportunity that accounts for holistic effects. The following chapter structures the methodology used to tackle the challenge, explicitly covering the reasoning used to elaborate it. It is separated in the four main blocks of the approach: Demand estimation, Simulation, Metamodel and Optimization. The results obtained are presented in Chapter 5 and, additionally, an experimentation with emergency shipments is introduced, made possible by the feature introduced in the Simulator. Lastly, conclusions are drawn, and the limitations of the present work are stated, along with the possibilities to enhance it.

Chapter 2

Theoretical background

This chapter aims at presenting a complete overview of the relevant concepts and proposed methodologies that steered the development of the project. The first section provides the reader with an understanding of the characteristics that distinguish the fashion industry and introduce causality to the methods presented for dealing with the associated complexity. Section 2.2 elaborates on the holistic supply chain definition and trends, with particular attention given to the elements with interdependence between logistic subsystems. In Section 2.3, demand estimation approaches are briefly introduced, as a data precedent for the explored analysis. Finally, in Section 2.4, simulation and optimization techniques are explored as well as hybrid methods that combine advantages of both, along with the presentation of published case studies on pertinent research questions.

2.1 Fashion Industry

Each industry's supply chain presents unique challenges that must be taken into account when establishing adequate strategic and tactical planning. Fashion & Apparel is no exception, and increasing competition has been shaping the way players operate. Christopher and Peck (1997) summarized four key defining characteristics: short product life-cycles, high volatility, low predictability and high amount of impulse buying.

2.1.1 Short life-cycles

Clothing items can be divided in basic (selling all year), seasonal (two or three seasons per year) and fashion, with shelf lives of 8–12 weeks (Nuttle et al., 2000). By definition, fashion items represent current, short-term demand, which means that in order to keep up with customers' expectations, stores must present new products frequently (Caro and Gallien, 2010), leading to the creation of multiple collections within one selling season (Christopher et al., 2004). Fast-fashion retailers introduce around 2000-4000 items per year (resulting in thousands of SKUs when we consider different color and size variations) (Nueno and Ghemawat, 2003). The inability to sell the product within its lifetime leads to markdowns at the end of the season (and the consequent reduction of profits) in order to clear the stores for new collections. This puts extra pressure to quickly and correctly allocate stocks to stores while maintaining limited inventories.

2.1.2 High Volatility

Through large scale surveys and interviews, Amed et al. (2018) consistently report volatility as the main challenge for fashion players, as purchase desires can be shifted based on multiple factors, such as weather or celebrity influences. For instance, Portuguese players reported lower than average sales during the Spring-Summer collection of 2018 due to uncharacteristically cold months that delayed consumer shopping trips for warmer weather clothing.

2.1.3 Low Predictability

Additional strain comes from the impossibility to correctly predict demand at the SKU level, due to the very low average sales per item per store per day. Although product group (e.g. dresses) demand is relatively stable, due to substitution effects, the personal nature of the products makes it hard to predict whether a specific pink t-shirt will sell (Collado and Martínez-de Albéniz, 2017). Extra complexity is added by the store dimension, which can have unique selling patterns and clients (Correa, 2007; Collado and Martínez-de Albéniz, 2017), as is typified by the case of holiday locations' stores whose demand skyrockets during the holiday season. Christopher and Peck (1997) state this uncertainty should be counteracted by short lead times, making the supply chain more agile to cope with demand fluctuations.

As consumer habits shape the demand curve, it is expected for seasonality effects to become clear when purchases are analyzed at the aggregate level. The peak found on the weekend puts pressure on the replenishment activities performed on the days after to fill shelves and bring stock levels to stable values. This effect has been accounted for in studies by Collado and Martínez-de Albéniz (2017). Moreover, special days and holidays are linked to an increase in customers as was proved by Nieves-Rodríguez et al. (2015), which concluded that Mother's Day and Christmas had a positive impact on sales in a Puerto Rican women's apparel retailer.

Forecast errors result in two main consequences (Mattila et al., 2002): stock-outs implying lost sales (by underestimation of demand) and end-of-season markdowns (by overestimation).

2.1.4 High impulse purchase

In order to maximize the fulfillment of demand, availability becomes of utmost importance to increase profits (Fernie and Grant, 2014), however, Corsten and Gruen (2003) reported worldwide overall retail out-of-stock rate of, approximately, 8%. According to empirical evidence collected by Zinn and Liu (2008), clothing stock-outs result in the majority of the cases in a lost sale, with 44% giving up on the purchase, 39% leaving and going to a competitor and only 17% delaying and coming back. These findings differ from previous studies by Corsten and Gruen (2003) in consumer goods, which, although in line for delayed purchases, recorded 9% probability of lost sales and substitution becoming the most common response. Apparel specific factors might be underlying the gap, with more frequent misplacement of inventory (for instance, in changing rooms), creating the need for reprocessing for placement on the shelf and lowering availability (Fernie and Grant, 2014). More recent research presented by Ovezmyradov and Kurata (2019) highlights the importance that omnichannel strategies have on active responses, as customers now have different options and, for instance, delaying the purchase might mean "order in-store, deliver home".

Collado and Martínez-de Albéniz (2017) analyzed the relationship between higher inventory levels (resulting in better presented displays) and higher sales, which were found to be positively correlated. However, new models of fast-fashion now leverage stockouts as a way to promote exclusivity and rush customers to stores (Fernie and Grant, 2014), leading to more frequent visits, as best evidenced by Zara, which declares 17 visits per customer per year on average, in comparison to three or four for non-fast-fashion competitors (Fraiman et al., 2010). This strategy is dependent on a responsive supply chain, capable of reading this customer surge and ordering quick production of more items. On the other hand, an availability mindset becomes unavoidable when dealing with fixed, pre-ordered quantities, which is elaborated on in the following section.

2.2 Supply Chain

2.2.1 Procurement

Cost reduction measures have created a trend for off-shore production in low-cost labor countries (Mattila et al., 2002; Nueno and Ghemawat, 2003), extending the supply chain and creating long lead-times between design and retail, which reduces agility and increases response times (Christopher et al., 2004). Collections need to be fixed months prior to the selling season and extra orders are not possible in useful time, which places unrealistic expectations on accurate long-term forecasts (Mattila et al., 2002).

As a consequence, and considering the specific challenges that are posed by the industry, new strategies have been shaping apparel supply chains, searching for better adaptability to demand and higher product assortment refreshment, as is presented in Figure 2.1. By separating basic products with stable demands, which can be manufactured in advance, from fashion, trending ones, fulfilled quicker by the use of local production, it is possible to obtain additional flexibility while keeping in mind cost concerns (Nueno and Ghemawat, 2003). Quick response (QR) techniques, leveraging information sharing between supply chain participants, to allow for quicker reactions to changes in demand and mid-season production, have been implemented in several industries with high volatility (Christopher et al., 2004). More recently though, Fast-Fashion concepts were born, and retailers redefined "seasons" by being able to quickly design, manufacture and put in stores new products throughout the year, with Zara becoming the first example of just how successful this strategy can be (Christopher et al., 2004). In an attempt to follow this trend, most players with less ability to restructure have opted for semi-planned management approaches (Martino, 2015), creating the illusion of constant redefinition of collections by manufacturing in advance multiple items in small batches, which they consecutively introduce in stores. Procurement activities, related to the relationships between suppliers and the company, are then simplified and the retailer becomes focused on the positioning (and corresponding timings) of its inventory. This type of strategy will be the core of this project and the one used by the case study presented, notwithstanding the possibility of implementation in other supply chains.

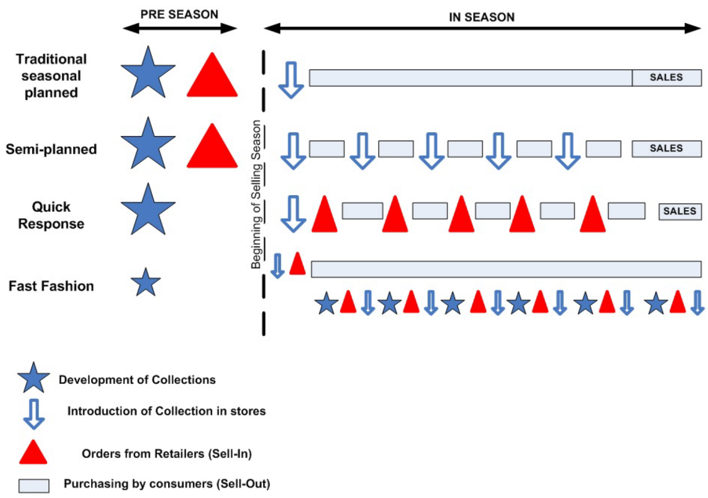

Figure 2.1: Fashion management approaches. Source: Martino (2015).

2.2.2 Distribution

Martins (2018) summarize the types of delivery modes that can be used to fulfill brick-and-mortar stores' demand from suppliers: warehousing, direct-shipment, cross-docking and milk-run. According to Kuhn and Sternbeck (2013), retail distribution via distribution center represents the clear majority of situations and is predicted to keep its hegemony. Aftab et al. (2018) refer as common practice for Zara competitors to ship products from distribution centers (where orders from suppliers are received) to regional warehouses and finally to stores, in contrast with direct shipment from a central warehouse in Zara's case. The last chain of these networks, where decisions of allocation and consolidation of orders are done, is the focus of this project. This postponement strategy optimizes the "time utility" by delaying shipments and the "place utility" by choosing the most adequate point-of-sale (Yang and Burns, 2003). Three main logistics subsystems are found (Kuhn and Sternbeck, 2013): distribution center (DC), transportation and store (see Figure 2.2).

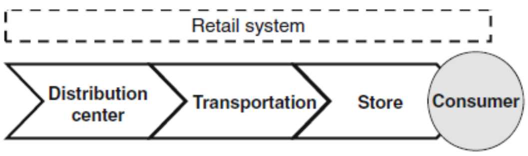

Figure 2.2: Retail logistics subsystems. Adapted from Kuhn and Sternbeck (2013).

For the first dimension, Rouwenhorst et al. (2000) divided warehousing processes in receiving, storing, orderpicking/consolidating and shipping. These are, in broad terms, not industry specific, although some products could require some additional handling, as, for instance, putting tags in clothes.

As for the store subsystem, Pires et al. (2017) devised a general definition of grocery retail stores based on the works of Kotzab and Teller (2005) and Reiner et al. (2013), presented on Figure 2.3.

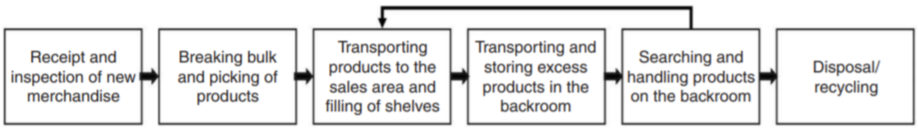

Figure 2.3: Breakdown of store operations. Source: Pires et al. (2017).

Some points of contact can clearly be found to apparel retail, as inventory is identically kept on shelves (or hangers) on the sales floor and in backroom spaces. This storage facility acts mainly as a buffer between deliveries and limited shelf space (Mou et al., 2018), taking on additional responsibilities on the case of fashion retail, such as strategic delay of presentation of new products or storing collections that had missing key sizes, a common practice in Zara (Correa, 2007). Ton et al. (2010) divide the in-store logistics of Zara in four processes: processing deliveries, managing product flows between backroom and selling floor, managing display areas and fitting rooms and performing physical audits of inventory. Although these can be fairly generic, Pettinger (2004) stresses the importance that targeted brand image and organizational goals have on the combination offered of self-service (involving preparing the merchandise), routine service (like receiving payments) and personal service (as giving advice and searching for sizes), as the three variations of customer service. This means that companies might aim at increasing the free time of store personnel to allow more interaction with customers or, on the other hand, increase task efficiency and personnel occupation.

Kuhn and Sternbeck (2013) identify five critical tactical elements that have inter-dependencies between logistic subsystems: store delivery pattern, store replenishment lead time, order packaging unit, store delivery arrival times and roll-cage sequencing and loading carriers. A deeper explanation based on their work is presented for the first two, keeping the others as exogenous variables, although further studies could implement improvements on those dimensions.

Delivery Patterns

Delivery patterns represent the combination of days (number and timing) in which a store receives inventory from distribution centers (Sternbeck and Kuhn, 2014), in a fixed weekly schedule. Receiving items steadily on the same days has been defended by researchers as a preferred practice, allowing for a fixed scheduling of the workforce, known routes and unloading conditions for trucks and overall operational stability (Holzapfel et al., 2016; Sternbeck and Kuhn, 2014; Gaur and Fisher, 2004). Moreover, it allows for periodic reviews of inventory at predefined times, in a consistent matter. Equidistant deliveries are also supported by Kuhn and Sternbeck (2013).

The frequency at which these deliveries need to happen is highly dependent on the type of products considered, as perishable goods or unpredictable demands will result in an increased number of shipments. Moreover, other store-related constraints might define frequencies, such as the size of the backroom, receiving capacity and sales volumes (Kuhn and Sternbeck, 2013). The major impact that differing order spacing will have is in order volumes, that is, if orders are placed less frequently, bigger quantities will be delivered each time. Martins (2018) mention that having larger orders allows for efficiency in transportation by reaching full-truckloads, however, this is hardly the case in fashion retail, as average weekly orders consolidate at around 3 pallets per store for the case study analyzed. Nonetheless, when considering outsourced transportation, it is advantageous to be able to maximize pallet occupancy, which might not be the case for daily deliveries to low rotation stores.

In terms of the specific days chosen, one effect that cannot be forgotten has to do with the seasonality of the demand. This translates into some days being better suited to avoid additional backroom storage and handling, as more items will be missing from the shelves at the beginning of the week, due to the cumulative effects of higher demand and, usually, no weekend deliveries (Kuhn and Sternbeck, 2013). The effects on the DC depend on the holistic integration of all the stores supplied, as an adequate intertwining will lead to smoother picking requirements and avoid requiring extra capacity (Sternbeck and Kuhn, 2014).

Replenishment Lead Time

The amount of time a store needs to submit an order before receiving the items represents flexibility that is shared between the store and the distribution center: if shorter, retailers are more capable of dealing with demand fluctuations and more reactive in orders. However, it comes at a price for the DC which must be able to prepare and ship items quicker, reducing opportunities for smoothing operations throughout the defined lead time. Transportation can also be affected by this just-in-time approach, reducing order consolidation opportunities or, depending on the contracts established with third-party logistics, become affected by higher tariffs (Kuhn and Sternbeck, 2013).

The possibility of emergency shipments (with consequently shorter lead times) under specific conditions has been studied in multiple variations in literature. In spite of it, most focus on the relationships between suppliers and company facilities, disregarding the entropy created and limiting the resulting effect to more expensive transportation or ordering costs (for instance, in Tagaras and Vlachos (2001)). Transshipments between stores are also evaluated as an alternative option to urgent deliveries in Chartniyom et al. (2007) and Liao et al. (2014).

2.2.3 Replenishment methods

Replenishment methods have been widely researched to find the optimal stock policy, with earlier publications dating back to the beginning of the 20th century (Bartmann and Beckmann, 1992). Although several variations can be mentioned, the ones tackling uncertain demand can be divided in fixed ordering quantity methods and periodic review methods (Waters, 1992). For the first approach, a constant amount of product is requested whenever a certain level is reached, assuming a continuous review of the inventory. On the other hand, periodic reviews assume checking stock at fixed time intervals, meaning that variable quantities will be ordered in response to the observed level. When considering multi-items (which is the case for a fashion retailer), a coordinated replenishment has been defended as cost-saving (Jayaraman and Tabucanon, 1984) and enabler of the definition of delivery patterns (Cardós and García-Sabater, 2006). Waters (1992) mentions how a periodic review can provide a simple routine of ordering, delivering and receiving inventory. This is usually done by ordering each R periods until an order up-to level S, which leads to the definition of a (R,S) inventory control system. Hybrid approaches combine advantages of time and quantity focus, which is the case of the (R,s,S) where an s level is established and only if stock levels are below s, will an order be placed for the cycle.

2.3 Demand estimation

As analyzed by Garro (2011), captured demand by sales data might not represent the full extent of real demand, if affected by stock-outs (Mou et al., 2018) or non-routine closure of stores (for instance, closing down for a holiday). The ability to correctly evaluate sales potential is extremely important for forecasting methods (Wecker, 1978), as the failure to do so could introduce an under-estimation bias in predictions, and, more relevantly to the matter at hand, for simulation purposes, as a way to detach historical data from the occurred scenario and allow for the testing of others.

Garro (2011) used the case study of Zara to test several options for determining weekly aggregated demand, from simple extrapolation from days with full demand fulfillment (that is, without stock-outs), to more complex combinations of inputs from size dependencies (of the same product) and weekly seasonality. Huang and Liu (2017) used a variation of one of the methods presented, creating a weighted average of the overall weekly sales weight of the days without stock-out. However, substitution effects are disregarded on these approaches. As fashion products' demand is highly correlated with availability, the fact that one product is not on the shelf might lead the customer towards buying other (for instance, by buying a green t-shirt instead of a blue one). Although methods have been presented for estimating these effects (Vulcano et al., 2012; Musalem et al., 2010), the complexity added to the methodology did not justify the approach over simplistic ones. Moreover, sets of substitutable products had to be defined, which can be highly debatable when talking about fashion products.

2.4 Simulation-Optimization

2.4.1 Optimization

Optimization techniques lean on mathematical programming to reach the best solution, as evaluated by objective functions, within the limits of problem constraints (Snyman and Wilke, 2018). They have been used on a multitude of applications and further extensions and adaptations are still being implemented.

In the fashion industry context, researchers have mostly focused on the allocation of fixed quantities of inventory from DCs to stores. Gallien et al. (2015) provided insights into the distribution of limited stock, deciding on first shipment quantities while maintaining enough warehouse stock for replenishment flexibility. This second phase of distribution was addressed by Caro and Gallien (2012), using an optimization model to increase sales. Both papers have used Zara as a case study, where shipment decisions are made simultaneously for all stores, aiming at increasing overall profits. More recently, Martino (2015) used the bees algorithm to solve the replenishment problem, assuming four fixed deliveries during the season.

Although extensive research has been published on the transportation optimization dimension, through attempts at solving the periodic vehicle routing problem, only a few articles were found that tried to find interdependence between the logistic subsystems of a retail supply chain for the problem of delivery patterns definition, which are presented below in chronological order.

A heuristic approach was used by Gaur and Fisher (2004), determining delivery routes, allocating trucks to them and readjusting departure times to balance DC workload. Cardós and García-Sabater (2006) analyzed the dependencies between inventory and transportation costs, defining ideal frequencies which are combined into delivery routes, in a search for cost minimization. Although the store dimension was considered and integration was introduced for transportation purposes, the authors failed to consider upstream logistics.

The first system to tackle interactions between the store, transportation and DC on this problem was developed by Ronen and Goodhart (2008), by incorporating limited backroom storage and DC constraints on the periodic vehicle routing problem of both dedicated and contract carriers. One important limitation of this research has to do with assuming equal distribution of demand throughout the week, neglecting seasonality effects and associated volume differences.

Kuhn and Sternbeck (2013) provided a theoretical framework for developing holistic models that encapsulate effects that certain decisions have on logistic subsystems of the supply chain. This allowed Sternbeck and Kuhn (2014) to present the closest approach to the one chosen for this research, minimizing overall supply chain costs in the definition of ideal delivery patterns for a network of stores. However, certain simplifications are considered, such as the assumption of full availability of products at the store, made possible by deterministic aggregate demands.

Further research by Holzapfel et al. (2016) included store bundling for transportation optimization when considering company owned trucks. This effect will not be taken into account in the model presented, as the case study analyzed outsources this service and possesses fixed destination-pallet costs.

Our methodology adds to the field, by including lost sales incurred by higher demand response times. This becomes of greater importance in the fashion context considered, as forecasts are extremely unreliable. The inability to correctly predict the relationship between delivery patterns and fulfillment of demand leads us to the use of simulation.

2.4.2 Simulation

Simulation has been extensively used as a tool to replicate the functioning of a real system, in order to obtain insights into how it operates or experiment with alternative scenarios (Shannon, 1975). To do so, a representative and simplified model is created, with its inherent relationships between agents relevant to the analysis. Banks et al. (2010) classified models into static or dynamic, deterministic or stochastic and discrete or continuous. A dynamic, deterministic and discrete simulation represents a system evolving through time, in a fixed relationship between inputs and outputs (without random influences) and in which variables change states in a set of moments in time. Through discrete-event system simulation, one can model such system, creating the state evolution through event triggers, while tracking individual items. Duong and Wood (2018) mention how this makes it the ideal methodology for inventory management simulation, where a set of rules will create specific effects for each product.

A decision support system based on a supply chain simulation tool is proposed in Longo and Mirabelli (2008), making clear the benefits that this approach has in understanding the effects of different demand pattern, inventory policies or lead times on a multi-echelon network.

Cerda and Monteros (1997) used a weekly simulation approach to experiment different c ("can order") parameters on a (R,s,Q,c) multi-item replenishment policy between retailer and warehouse for a cardboard box marketing firm with lot size Q. Although this case study presented different goals and tested variables, the main events are in their core similar to the present research objective, as inventory levels, arrivals of products, sales and requests to the warehouse are computed for each period.

In the fashion industry, simulation has been mostly used to model store dynamics and customer responses, which is the case of De Marco et al. (2012), where RFID technology was analyzed through system dynamics, and Al-Zubaidi and Tyler (2004), who modeled client behaviour to evaluate the impact of different QR replenishment policies on lost sales. On a different subject, Garro (2011) provided a simulation approach to experiment with new allocation definitions at Zara and Iannone et al. (2013) used simulation to introduce a push-pull strategy to product allocation based on forecasts. However, the simplistic shipment policy adopted in both cases (once at the beginning of every period) turned the design quite straightforward, although demand estimation was implemented on the first one.

2.4.3 Metamodel-Based Simulation Optimization

The logical consequence of the identification of the benefits of simulation and optimization approaches separately is to leverage strengths in new methodologies, reaching the fine detail of simulation for an optimized solution. Problems with interrelated decisions and complex causal effects are hard-to-model but can be dealt with in simulation without much mathematical sophistication (Figueira and Almada-Lobo, 2014). Although simulation experiments can try out different combinations of inputs to reach insights into the detailed results they provide, in numerous cases, the extensive enumeration of all possibilities is computationally prohibitive (Carson and Maria, 1997) or even impossible due to infinite search space. Optimization techniques can tackle this shortcoming by limiting scenarios and providing solutions in useful time.

Figueira and Almada-Lobo (2014) performed an extensive review and classification of hybrid Simulation-Optimization (S-O) methods, distinguishing between Solution Evaluation, Analytical Model Enhancement and Solution Generation approaches.

One particularly relevant method to overcome expensive simulation tests is the use of metamodels (Wang and Shan, 2007), defined as models of models (Jalal et al., 2013). After running the simulation for several points, an approximated and simplified relationship between inputs and outputs can be defined, which is used for deterministic optimization (Figueira and Almada-Lobo, 2014). Kleijnen (2005) proposed a 10-step methodology for correctly implementing metamodels. Related to validation, coefficient of determination (R^2) and cross-validation are presented as alternatives. Despite the immensity of options which deal with specific challenges, April et al. (2003) state that standard linear regression is still the most commonly used technique.

There is not extensive published research incorporating simulation-optimization methods on case studies of retail inventory problems. Duong and Wood (2018) focused on the benefits of simulation for finding consequences of chosen inventory policies for perishable products when used in a recursive manner with optimization techniques. A case study using this methodology is presented by Myers (2009). Moreover, Hachicha et al. (2013) used a metamodel-based optimization simulation approach to define ideal values of s and S in a (s,S) continuous replenishment policy. Using a factorial Design of Experiment, they perform simulation runs for different combinations of factor levels whose estimated relation with the observed response is used for optimization. Although not completely related to the problem analyzed, the mentioned examples provide insights into the potential of the methodology.

Chapter 3

Problem description

As has become clear from the presented research, the definition of delivery patterns has to account for the different consequences originated throughout the logistics subsystems. In order to do so, the fashion retailer's relevant characteristics are presented in Section 3.1, including operational details. Section 3.2 builds on this to define the challenge that is to be tackled, as well as the shortcomings of the present approach used by the retailer.

3.1 Case study context

The fashion retailer that motivated the case study operates a centralized distribution center that supplies around 150 company-owned stores in Portugal and Spain, with a total of 9 business units (ranging from Baby Apparel to Adult Non-apparel). Using a semi-planned management approach, where designs are completed months before the selling season and manufactured in low-cost countries, the company manages to mimic fast-fashion practices by constantly introducing new collections in stores. Each combination of color-size product is defined as a Stock Keeping Unit (SKU), possessing a unique identifier. Every week, close to half of the SKUs shipped to stores constitute first shipments of new products, not unlike what is practiced by Zara (Gallien et al., 2015). This fact provides additional complexity to the problem but at the same time more flexibility to smooth operations.

Assuming upstream processes of procurement as exogenous to the analysis and consequently fixed quantities of each product, the three relevant logistics subsystems are store, transportation and distribution center (DC). An in-depth breakdown of the retailer processes and constraints for each dimension creates awareness of their holistic integration and is presented on the next subsections. It is important to mention that in order to provide a general approach to the problem, certain specific details and practices were simplified when, although implemented, they do not significantly alter the methodology or the results obtained.

3.1.1 Store

The last chain of the supply chain is where customer demand is realized. Mainly, store personnel is focused on adequately presenting products (with the former processes it entails) and giving attention to the client. Following this customer-oriented principle, this fashion retailer's staff allocation is usually established with some redundancies, allowing better interactions with customers, when necessary.

Three main in-store processes were identified and detailed below: processing deliveries, selling and ordering.

Processing deliveries

The products arrive at their destination in the designated delivery days (from Monday until Friday) inside boxes in pallets, after replenishment orders issued by the store or new product introductions governed by centralized decision-makers. The latter case mostly occurs at the beginning of the week (Monday or Tuesday), as it is fixed that, since the demand peak of the weekend is burdened by the absence of deliveries, all clothing items must be pushed to refill shelves.

The personnel retrieve the pallets from the predefined location (in some cases in front of the entrance of the store, others from a centralized unloading area), and remove the protection film wrapping them. Afterwards, boxes are checked into the system one by one, and the shipment is verified. Clerks open and remove the products, while separating them by business unit and by the type of products (new item or replenishment). After preparation is completed (removing plastics, applying tags or putting in hangers), products are either carried to the sales floor or left in the backroom for when needed. On average, around 20% of the stock is kept in this temporary storage, consisting mostly of new products that do not yet fit with appropriate visibility on the racks. Depending on the delivery time, most of this process can happen while the store is still closed, allowing full dedication to it.

When shipping from the DC, some capacity constraints need to be accounted for, since the area for receiving pallets and backroom storage capacity is limited and different among stores.

Selling

In customer-oriented fashion retail stores, as is the case, employees are expected to dedicate a fairly large part of their time to direct contact with clients facilitating purchases, be it by giving advice or helping find a missing size. When demand is fulfilled, the clerk will realize the transaction, receiving payment and recording the sale. In this high impulse buying context, availability becomes extremely important, as stock-outs will (to some extent) create lost sales. Inability to find a size or color for the desired product might create a cost to the retailer by missing out on the opportunity for a full-priced purchase and being obliged to use markdowns to clear inventory at the end of the season.

Ordering

Depending on the delivery pattern and on the lead time, automatic systems will place orders, accounting for inventory position (on-hand added to on-order) and an order up-to level (OUTL). This periodic review system (in contrast to a continuous option) allows for the consolidation of shipments in the defined delivery days and is known as a (R,S) replenishment policy, where every R periods, inventory will be ordered to bring stock up to S units.

$$OUTL = \lceil \mu \times (LT + RP)_{weighted} + \sigma \rceil + 1 \tag{3.1}$$

The generic equation 3.1 defines the OUTL, which is specific per SKU and per store. The average daily sales of (up to) the last 4 weeks (μ) is computed and multiplied by the lead time (LT) and review period (RP). This store-fixed value is calculated considering the weekly seasonality presented in Table 3.1, which weights each day's time until replenishment: if a store orders on Tuesday to receive on Thursday, three (days) is multiplied by Monday's 12%, accounting for the time this share of the sales take to be restocked. Adding the standard deviation of the period (σ) provides safety stock for variations of demand. The resulting amount is rounded-up and 1 unit is added to define the maximum level. Every Monday, these parameters are recalculated to include newly available data. Accounting for the large number of SKUs carried, and the consequent low daily average sales per each one, usually only 2 pieces are in stock for most items.

Table 3.1: Weekly seasonality of global chain sales for 2018.

Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
12%	11%	12%	13%	15%	19%	16%

3.1.2 Transportation

The analyzed retailer outsources all transportation to a third-party logistics provider (3PL), which presents benefits and disadvantages worth considering. On one hand, there is a considerable simplification to the distribution process as all routing optimization decisions are out of scope. Limitations on capacity are also disregarded, as the service provider operates a large fleet of vehicles for multiple companies, being able to cope with variations on shipment quantities. However, the enabled flexibility comes at a price, since the cost of shipping a pallet becomes higher and fixed per destination, putting extra pressure on guaranteeing efficient deliveries in terms of timing (reaching the store when most needed) and pallet consolidation (increasing volume per pallet transported by combining replenishment needs from several days). Moreover, since orders need to be consolidated with other shipments of the 3PL to reach competitive pricing, longer lead times are imposed, meaning that for a product to arrive at a Portuguese store in the morning, it must leave the distribution center in the evening of the previous day. Spanish stores are also included in the analysis with one extra distribution day.

3.1.3 Distribution center

The centralized distribution center is responsible for the processing and allocation of the inventory received from manufacturers, being operational only from Monday until Friday. After initial reception activities, the items are organized and placed on a designated spot. Every day, shipment orders are consolidated, and picking takes place, collecting the total units from each SKU necessary to fulfill all store requests. Afterwards, a product sorting machine (sorter) is fed with these products and begins distributing them into boxes, which will only contain items going to the same store. This process constitutes the bottleneck of the operation, limiting the number of units dispatched from the DC to the pieces it can separate multiplied by the time it is operating. Afterwards, boxes are prepared for shipment by being placed on a pallet and wrapped with protection film, getting loaded into the 3PL trucks.

The analyzed DC operates large volumes per day for multiple end retailers with limited capacity, which means that, in order to efficiently use resources, setup times must be minimized and picking consolidated, allowing the sorter to run full courses, usually from the beginning of the afternoon until midnight. Considering the transportation shipping time (starting from around 5 p.m.), loads will only be shipped the next day, accounting for the usual 48 hours of lead time for national stores.

3.1.4 Costs overview

In order to provide a clear definition of the costs for the methodology to be presented, Table 3.2 summarizes each activity with the corresponding measuring unit. To allow generalization, these have been grouped into broad definitions.

Table 3.2: Costs per subsystem

Subsystem	Dimension	Units
Store	Receiving	Pallet
	Processing	Item
	Lost sales	Item
Transportation	Delivery	Pallet
Distribution center	Picking	Item
	Sorter	Item
	Shipping	Pallet

3.1.5 Seasonality

Apart from the weekly seasonality already mentioned, it is quite straightforward to understand that different periods will have different demand patterns throughout the year. Figure 3.1 depicts the landscape of sales quantities (in units) and profits (selling price minus acquisition costs) per average store (accounting for opening and closing of spaces) for each week of the year 2018, for the retailer under analysis. The values are presented as the weekly percentage over the year total, and a dotted line is included to represent the average value if all weeks had the same weight.

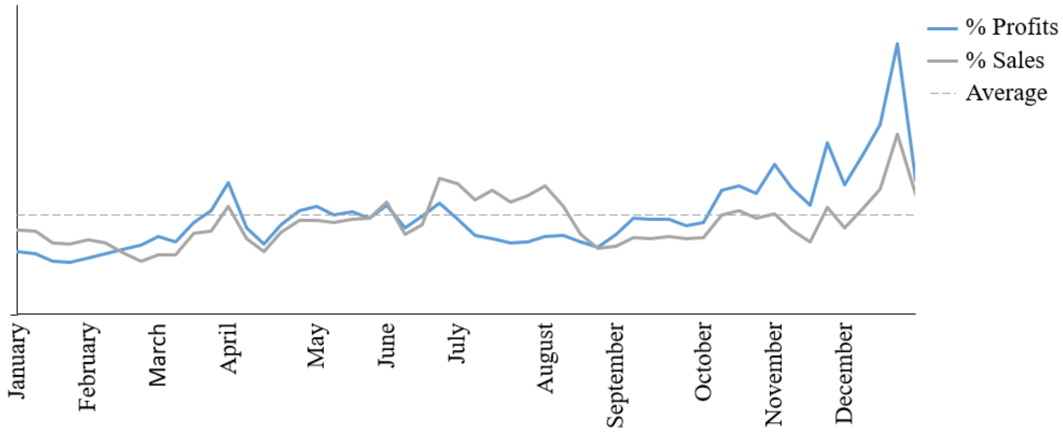

Figure 3.1: Sales quantities and profits per week for the year 2018.

Although sales volumes show higher volatility, profit distinctively peaks in April (explained by Easter) and at the end of the year, which includes Christmas and Black Friday. One effect which can be observed is how markdown periods (typically after Christmas through February and during July and August) result on the lines crossing, representing that, though larger volumes are being sold, the lower prices practiced lead to lower profits. In other words, although markdown periods are strong in units purchased, they are less relevant for global profit. On the other hand, the months from October on show a clear trend towards high returns, and above average importance. This long-period seasonal consumer demand might create the need for specific delivery policies, as higher transportation costs can be offset by higher profit margins.

3.2 The challenge

Keeping operational practices stable, the challenge undertaken aims at establishing the ideal delivery days per store, and consequent delivery frequency, while accounting for the impacts in all supply chain dimensions. Having described the main activities involved in the replenishment process, one can now portray the definition of individual delivery patterns as the ideal trade-off between the service level provided and related fulfillment costs.

On one hand, should products be shipped less frequently, pallet occupation can be maximized, which will, in most cases, diminish the total number of weekly pallets, decreasing expedition costs at the DC and avoiding trips that deliver only a few boxes. Conversely, a higher delivery frequency (with added benefits if the specific days chosen are optimized) means that sales can be replenished faster and turn stock-outs less frequent, which can lead to increased profits if the item ends up in markdowns or does not sell in stores.

Most store operations are dependent on the number of units delivered, but not on whether they are shipped together. Moreover, for the case study at hand, we are considering a certain idleness percentage by the store personnel, which is capable of absorbing variations in the number of delivery days. This assumption is an accepted belief by company stakeholders, but could easily be considered in the model for different conditions. Regarding store space constraints, a higher delivery frequency can be imposed.

The problem gains an additional degree of complexity when taking into account operational restrictions at the DC, as ideal scenarios per store must be balanced to find a global optimum. Smooth sorter utilization (as the bottleneck of the operation) requires that stores are efficiently intertwined in the specific days chosen for shipment. This process must also consider new product deliveries, which provide flexibility as they can be distributed throughout the planned introduction week (while complying with business rules).

Since stability is valued for routine operational details, the delivery patterns introduced should be stable for several weeks. However, as holiday season variations were identified, different policies for this period should be evaluated, splitting the year in two.

It is important to mention that inventory costs are not considered (unlike the methodology presented by Sternbeck and Kuhn (2014)) since they are deemed as fixed and unrelated to the timing and placement of the previously bought products.

3.2.1 Current approach

Until this moment, the retailer relied on years of experience to perform an *ad hoc* approach to solve this problem. Store delivery frequency was mostly dependent on sales volume, where 'best' stores were assigned with daily deliveries, and adjustments took place whenever peak periods were expected or stores expressed lack of storing space. This accounted for an average of 3.3 weekly deliveries per store. When comparing to Zara's twice per week global shipments (Nueno and Ghemawat, 2003), an improvement opportunity was found, as this fast-fashion giant has proved to be a role model of industry's good practices. Furthermore, the distribution center often reported inability to comply with requests, as capacity constraints were exceeded, leading to close to 10% of units leaving later than they were scheduled to. This provided additional motivation for better balancing and synchronization through intertwined delivery patterns.

Chapter 4

Methodology

The present chapter details the main steps of the methodology undertaken to overcome the problem of finding fitting delivery patterns for the challenges posed by a fashion supply chain. Invoking the categorization proposed by Jalali and Nieuwenhuyse (2015), the problem can be described as single-echelon, multi-item and finite horizon with deterministic lead time and an (R,S) replenishment policy. Following a summary of the proposed approach disclosed in Section 4.1, the demand estimation technique is explained. An in-depth description of the simulation model is provided afterwards, as well as of the validation procedure. In order to reduce the need for simulation runs, a metamodel is defined in Section 4.4, ending with the optimization model implemented.

4.1 Proposed approach

The proposed approach to address the definition of delivery patterns is summarized in Figure 4.1.

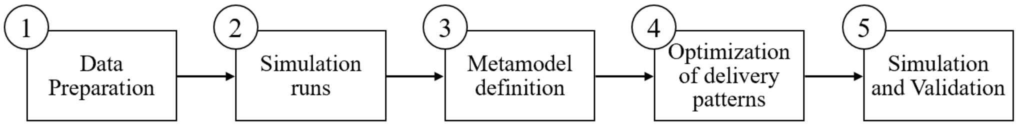

Figure 4.1: Main steps of the methodology presented.

The first phase encapsulates all steps taken towards gathering unbiased data for simulation. Historical sales, transfers and stocks were provided for each store, which were extremely dependent on allocation decisions and timings. In order to keep strategic store-specific policies fixed, each product's lifetime at each store was traced, setting the date for the first shipment and only allowing sales until the last recorded purchase. However, as stock-outs accounted for around 5% of total day-SKU-store entries, there was the need to estimate demand for this information blanks.

For the second step, a simulator was constructed, capable of reading data inputs and, following business rules in regard to shipments and SKU parameters, recreating sales, stocks and transfers from the DC, while accounting for fixed quantities of total products. Variable delivery patterns were fed, creating results for several scenarios tested. The use of simulation allows perceptions on how lost sales at the beginning of the collection can result in lower profits by higher markdown quantities.

As simulation runs are computationally expensive and combinations of possibilities extensive, a metamodel-based approach was used to estimate store profits for differing replenishment speeds. Using simulation outputs, an average week was defined for the two periods considered (until September and from then on), which associated weighted replenishment lead time to each delivery pattern. Through linear regression, this independent variable determined profit per store.

The steps undertaken so far provide all inputs needed to perform optimization, considering the cost-benefit trade-off, incorporating store and DC capacity constraints and new product shipments with higher flexibility in the delivery day. After enumerating all the combinations of store-delivery pattern, a holistic solution is found maximizing overall profit, for the time frame of one week.

The last phase of the methodology tests the configuration found, validating constraints and finding detailed costs for an entire year of operation. An additional optimization module was introduced beforehand to distribute new product shipments throughout the week, the need for which will become clear after the definition of the simulator in Section 4.3. These last steps will not be detailed further due to the simplicity of the approaches used.

4.2 Demand estimation

One of the inputs needed for simulation is a complete assessment of the demand for each combination SKU-day-store. Using historical sales data for this end is possible if all potential purchases were fulfilled, which is not the case in days in which a stock-out occurred.

First of all, for the analysis at hand, a stock-out is identified only in days that fall between the first and last sale of a specific item in a specific store, keeping strategic timing decisions fixed. Moreover, it was observed that for most days ($\sim$99%) in which inventory is present at the store, at most one unit is sold, which goes in line with the low stock kept for each SKU. This led to the definition that only occurrences where no products were stored at the beginning of the day should be corrected for demand. This specific detail means that when returns to the store happen, resulting in a positive stock at the end of the day, that day should also be evaluated for potential sales. An added reason for this is the common practice of not processing a product back to the sales floor immediately afterwards.

The core of the methodology used to build the model is based on the one presented by Huang and Liu (2017), extended to incorporate specific characteristics of the problem at hand. Table 4.1 summarizes all the following used notation.

In order to include weekly seasonality, the first step is the definition of the weight each day $d \in D$ carries on global weekly sales, presented on equation 4.1. Since we are considering nine different business units $u \in U$, with possibly distinct sales patterns, the weights were defined for each one. Moreover, to account for a whole year of data, this calculation was done per week, catching non-regular volumes, caused by factors such as mid-week holidays. W_{udw} then defines

Table 4.1: Table of notation for demand estimation.

Sets and indices	
$m \in M$	SKUs
$u \in U$	Business units
$m \in M_u$	Subset of M for the SKUs that belong to business unit u
$d \in D$	Days = {Monday, Tuesday, Wednesday, Thursday, Friday, Saturday, Sunday}
$w \in W$	Weeks
$s \in S$	Stores
Parameters	
S_{udw}	Sales quantity of business unit u on day d of week w
S^m_{dws}	Sales quantity of product m on day d of week w for store s
V^m_{dws}	Binary variable that indicates whether product m was not on stock-out on the beginning of day d of week w for store s
Variables	
W_{udw}	Weight on sales of business unit u on day d for week w
D^m_{ws}	Demand for product m on week w for store s
D^m_w	Demand for product m on week w
SF_{uws}	Store factor for business unit u on week w for store s

the weight of day d for business unit u on week w, by dividing the sales quantity of u on day d of week w, S_{udw}, by the sum of the whole week.

$$W_{udw} = \frac{S_{udw}}{\sum_{d=1}^{7} S_{udw}} \qquad u \in U, d \in D, w \in W \tag{4.1}$$

Afterwards, estimated demand of each SKU ($m \in M$) was found for each week $w \in W$, using the relation in equation 4.2 (which is based on the notation of Huang and Liu (2017)), by adding, for all stores, the sales for the days in which no stock-out occurred and dividing the result by the sum of the weights those days possess for the specific business unit u to which the m product belongs. A binary variable, V^m_{dws}, was introduced to represent each day with sales possibility (that is, when there was inventory at the opening of the store).

$$D^m_w = \sum_{s \in S} D^m_{ws} = \sum_{s \in S} \frac{\sum_{d=1}^{7} S^m_{dws}}{\sum_{d=1}^{7} W_{udw} \cdot V^m_{dws}} \qquad m \in M_u, w \in W \tag{4.2}$$

The store dimension, SF_{uws}, is introduced at u level (reducing instability of low selling SKUs), by creating a factor that represents how much the sales volume for each specific s store represents to the global amount, in a certain week (summarized in equation 4.3). It is calculated by accounting for full demand at the store level for that w and u, divided by global demand for the same week and business unit. The definition of this variable allows capturing different store rankings throughout the year, caused, for instance, by increased influx to stores by the coast during the Summer or specific Portuguese or Spanish holidays that result in country-located peaks.

$$SF_{uws} = \frac{\sum_{m \in u} D^m_{ws}}{\sum_{m \in u} D^m_w} \qquad u \in U, w \in W, s \in S \tag{4.3}$$

As introduced in Chapter 2, fashion items are highly availability-dependent, meaning that if a specific item is not on the shelf, the customer will (with a certain probability) replace it with an adequate substitute or, to some extent, delay the purchase and wait for stock to be refilled. That said, an additional step was incorporated to avoid overestimating demand when another item was purchased instead, by multiplying the average demand calculated by the estimated probability of actual lost sale. According to Zinn and Liu (2008), quitting or going to a competitor after a clothing stock-out happens for 82% of cases (with "delay" closing the gap to 100%). However, the analysis disregards substitution for other products inside the store, as the behaviour is measured after the client exits. Average European consumer responses (for several retail categories) found by Corsten and Gruen (2003) was 36% for the same behaviours, 17% for delayed purchases and the remaining substituting the stocked-out item. As company stakeholders, for the case study considered, believed that customers' trips were too infrequent for delaying the purchase, a value of 53% was implemented. This approach is meant to avoid altering selling quantities of other SKUs that were sold as a substitution but that are much harder to identify. The final formulation is stated in Equation 4.4, finding the demand for each m product, in day d of week w and store s.

$$D^m_{dws} = D^m_w \cdot W_{udw} \cdot SF_{uws} \cdot 0.53 \qquad m \in M_u, d \in D, w \in W, s \in S \tag{4.4}$$

This decimal value was discretized using a Poisson distribution and fixed for each entry, avoiding a stochastic simulation since the additional benefits did not outweigh its complexity.

The global results found for demand represented a 5.5% increase in comparison to recorded sales, which was considered adequate taking into account the out-of-stock total percentages and the business insights of company stakeholders involved in the project.

4.3 Simulation

The following section focuses on the specification of the logic incorporated for simulation. Through exogenous DC inventory quantities, the model takes as input initial conditions and is able to perform replenishment to stores for multiple scenarios. It is prepared to deal with multiple delivery patterns throughout the year and allows activation of 'emergency' quicker shipments for stock-out items. Detailed, SKU-level results are obtained, allowing experimentation of different measures and understanding of their consequences.

4.3.1 Model

The model was built using the general-purpose C# programming language, through the Visual Studio software. The deterministic nature of the simulation allowed this flexible approach to become a suitable option, fully replicating practices and enabling complete creation from scratch. Interactions with databases were incorporated, in order to read data and output solutions.

In order to turn the sequencing more understandable, the lines of code that reproduced the system can be illustrated using the blocks presented in Figure 4.2.

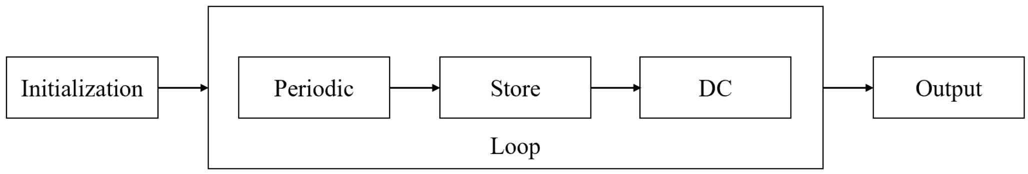

Figure 4.2: Main blocks of the simulation model.

In a macroscopic overview, after user inputs are introduced, the program starts by connecting to the database and initializing the variables, after which it enters a loop that runs for each day of the year (starting in February to allow some historical data). This loop encloses periodic, store and DC operations, exiting for a data output code block. Each relevant process will be detailed further.

Initialization

Before the simulation itself can begin, parameters need to be fixed for the run. Firstly, the number of different simulations to be performed is defined, in order to eliminate the need for waiting to introduce new parameters to be tested. Next, for each scenario, the user specifies whether to split the year in two (following the year seasonality identified in 3.1.5), and which set of delivery patterns to use (from the ones previously available in the database). One interesting feature implemented in the simulator is the possibility to override these delivery days for stock-out items, creating an emergency shipment to the store, with a shorter lead time by one day (that is, Portuguese stores are supplied in 24h, and Spanish ones in 48h). The activation of this policy is established also by user input.

Afterwards, code modules communicate with the database to establish:

- Store data: for each store, initial stock is defined for each SKU, along with the timing each product is to be on the shelf (first and last date);
- DC data: initial inventory (and arrivals from suppliers) and first shipment time and quantities;
- Client data: demand per store per SKU per day and store returns;
- Historical data: sales of the previous month and previous requests for shipments (for the first days of simulation).

This information is stored in dictionaries, relating unique keys with the correspondent values. These variables can be then accessed throughout the code for reading and writing purposes. Two cases where relevance and complexity are distinguishable require a more elaborate description.

The 'Timings' dictionary establishes when a certain SKU is available for selling at a certain store, fixing the first and last date, that is, <Sku,Store> $\rightarrow$ <FirstDate,LastDate>. One clear disadvantage of this rigid replication (which is necessary to avoid estimating demand for other periods/stores) is the fact that misalignments with allocation decisions taken in terms of quantities can seize products in stores that will never be able to sell them. Since we are dealing with fixed overall inventories, this will reduce global sales. In order to limit this effect, an additional steering variable was introduced in the model: the 'MaxDemand' dictionary. This acts as a ceiling for quantities to be requested by stores at each point in time, being defined at the beginning of the simulation for each SKU-Store (with the equivalent of a crystal ball) and decreased as demand is evaluated each day. It is easy to understand that replenishment flows will become much more fitted to historical ones, although first shipments could, nonetheless, introduce unsalable inventory.

Due to data availability limitations, the simulation was performed from the first of February, in order to include 4 weeks of sales transactions to calculate OUTLs, which is done for the first time in this initialization phase. Moreover, as DC inventory is to be shared with other channels (wholesale and other markets), transactions made to stores within scope were used as a proxy for the total quantity available to be shipped so that availability overestimation is avoided. In order to capture shipments from suppliers for permanent products (which can happen throughout the year), this method only considered data for one month ahead.

Periodic operations

Upon entering the simulation loop, the day variable is incremented to represent the present day, followed by the definition of the day of the week it refers to and of the demand per item per store (which is read from the database). Afterwards, several condition-based methods can be executed.

Firstly, as already mentioned, at the beginning of each week (that is, if the day represents a Monday), DC stocks are updated to include new historical transactions for one extra week, accounting for a full month. Furthermore, OUTL need to be redefined to include new sales transactions, representing potential increases or decreases on target inventory levels. This process is also done every Monday, replicating the policy introduced in 3.1.1, as reproduced in Algorithm 1. In order to provide stability to the parameter, a week is taken as a trial period where no changes are made to the predefined quantities shipped on the first introduction.

Additionally, if two yearly partitions are considered for delivery patterns, a function is activated on the day defined for the change, reading new shipment days per store from the database.

Store operations

The store dimension possesses a central role in the operations performed in the simulation, as most outputs to be observed depend directly on the interactions happening at this stage. After

Algorithm 1: Pseudo-Code for calculation of OUTL for each SKU-Store combination.

```
foreach timing in dictionaryTimings do
    if first shipment happened more than 1 week ago then
        total = Min{28;days on sale}
        aux = 0
        while aux < total do
            aux = aux + 1
            check sales for (total - aux) days ago
        end
        if there were sales then
            calculates OUTL
        else
            OUTL = 1
        end
    else
        OUTL = First shipment quantity
    end
end
```

initialization, each store is already paired with SKUs, identifying timings for each, and initial stocks are stored in a dictionary variable for each SKU, as well as the parameters for the OUTL.

Every weekday, the first step is updating inventory to include shipments that arrive on that day. For each one, the dictionaryStocks is either updated or a new entry is included, if that SKU was not yet present in the store. Additionally, the request that originated the delivery is eliminated. "Virtual" shipments are scheduled for Sundays for new introductions, so the same process is repeated on this day, with the exception of finding a request (the reasoning for this is explained in 'DC operations'). Afterwards, variables containing shipments data are dumped to an external file and cleared.

Next, in-store processes need to be replicated through logical sequences. Since OUTL are reevaluated every week and ordering can take place whenever there is a request window (or every weekday for emergency shipments), the obvious methodology would be to iterate through every combination of SKU-store to find if it was relevant for purchases or ordering. Exactly to avoid this time-consuming process (since fictional associations would be extremely common), the dictionaryTimings is used for looping purposes. This simplifies the computational effort, by limiting the field to pairs that are or will be pertinent. That said, every time an evaluated entry is on its last day, it is flagged for elimination. Algorithm 2 summarizes the activities performed for each combination.

After assessing that the SKU-store duo is active, the code starts by reading stock, demand, returns, maxDemand and the OUTL. MaxDemand is updated to reduce the demand of the present day and the ordering processes begin, to mimic the fact that this is usually done at the beginning of the day before clients enter the store.

The request creation is fulfilled after specific conditions tackling OUTL, maxDemand and

Algorithm 2: Pseudo-Code for Store operations per timing.

```
foreach timing in dictionaryTimings do
    if current day is between timing interval then
        reads stock
        reads OUTL
        reads demand
        update maxDemand
        if conditions are met then
            creates request
        end
        if stock > 0 then
            sells
        end
        performs returns
        updates stock
    end
end
```

Timing dates are checked, depending on the lead time type examined (a summary is provided in Table 4.2). Firstly, two shared filters are applied: it cannot be the last day of the timing (as, otherwise, stock would arrive after it is useless) and maxDemand has to be positive. After that, the two shipping modes are evaluated, starting by assessing the need for the quicker one, conditioned by whether it was configured for the current simulation run. Emergency requests are only permitted on weekdays (to always allow 24h for preparation) and if inventory position is equal to zero, which requires checking requests for on-order quantities. Moreover, the maxDemand should be larger than inventory position minus demand for the day to justify ordering more of that product for that store, in an allocation sense. At this stage, a new concept was introduced: considering that regular replenishment happens at least once per week, there was not a need to order aiming for maximum inventory, but only for enough to wait until the next scheduled delivery. After preliminary testing of more complex alternatives, 30% of the OUTL was established as a good trade-off between amounts transferred and sales potential, as recommended by company stakeholders. An additional limit is imposed before ordering to guarantee that the quantity is at most equal to the maxDemand, after which DC stocks are checked and reduced (if possible) to create a request. The request is stored with the date and quantity, as well as two boolean variables that state the urgency and whether it has already been shipped. An identical process is implemented to generate regular requests on 'request days' (established with sufficient lead time from the delivery day), being worth mentioning that on-order quantities are checked again to avoid double requests, when not needed. One important feature that is implemented is the ability that emergency requests have to 'steal' inventory from non-urgent ones from the same day, whenever the inventory at the DC is not enough to fulfill all orders.

After the completion of the ordering process, client interactions begin. Demand is compared with store stocks and the maximum value between the two is recorded as a sale. Afterwards, returns from that day are added to available inventory, in line with the practice of not processing

Table 4.2: Conditions for each type of shipment request.

	Regular	Emergency
When	Request window	Weekday & Activated
	Not last day of timing	
Conditions	OUTL > inventory position	Inventory position = 0
	maxDemand > inventory position - demand	
	maxDemand > 0	
Quantity	OUTL - inventory position	OUTL $\times$ 30%

the inventory beforehand. A final stock update is performed and the loop continues to the next SKU-store combination.

DC operations

The DC code processes requests from stores, creating shipments to be received by them, acting, essentially, as a delay module. Each weekday, unfulfilled requests are checked for shipping possibility, which means, both lead time compliance (at least two weekdays for Spain and one for Portugal) and delivery window existence (for that store) on the next weekday. Depending on whether emergency shipments were allowed, an attribute can identify these cases, which will be verified only for a more recent request date. If conditions are met, a shipment will be added (see Algorithm 3).

Algorithm 3: Pseudo-Code for DC Operations.

```
if weekday then
    foreach request in dictionaryRequests do
        if request not yet fulfilled then
            if emergency shipment then
                check if request was made at least yesterday for Spanish stores
            else
                check if request has predefined lead time and if the store has a delivery
                  window on the next weekday
            end
            if valid check then
                add shipment with appropriate flag
            end
        end
    end
end
```

At the beginning of the run, first shipment quantities per store per SKU are fixed to replicate strategic decisions, meaning available inventory must be limited accordingly. Since demand series were fixed to historical timings, it is important to guarantee that products reach the stores at least on the day they did on reality, so as not to miss the first days of purchases. However, we aim to use the flexibility of setting the specific date of the shipment for new introductions to balance DC workload over replenishment necessities. The resulting approach leads to the "virtual" distribution of inventory on the Sunday before the introduction week (even though sales will only be possible from the first day of the timing for each SKU-store) and an additional optimization engine is used after simulation to allocate them to the "real" shipment. This means that shipments might be happening after simulated sales, but it was considered that they would simply be delayed in reality (instead of Tuesday, happening Wednesday, for instance), not affecting global profits.

Output

After the simulation has run until the end of the year, the calculated data is sent into appropriate databases for analysis. This step includes writing sales records from the temporary variable used so far and uploading the files with shipments data.

Additionally, the CPLEX optimization module is run to distribute first shipments appropriately throughout the week and final results are exported. This last phase is performed merely for validation purposes so capacity limitations can be identified and shipment costs calculated correctly.

4.3.2 Validation

As defended by Banks et al. (2010), the validation of a simulation model is of utmost importance for increasing confidence in its ability to closely reproduce the system's behaviour, enabling experimentation, and for increasing credibility on the recommendations originated.

For the model produced, two main outputs are generated after the run: sales and shipments. Since we aim to get higher profits and more consolidated deliveries after experimentation, it should be guaranteed that by using historical delivery patterns we obtain an approximation of historical sales and pallets transported.

For the first dimension, a slight overall reduction of 2.5% was found when comparing units sold, and less 2.2% regarding profits. In order to investigate the underlying reasons for the difference, the simulated sold quantities per week were plotted against expected values, which is depicted in Figure 4.3.

A quick look at the graph immediately shows that the main variation seems to be located during the summer markdowns (that is, around July and August) and Christmas season, where quantities peak. In fact, removing these three months from the calculation lowers the deviation to 1.3%, for the remaining eight. However, after carefully analyzing the data provided, one can account most of this effect to errors in records during these chaotic periods of inventory push. Since not all leftover stock shipped to stores for end-of-season clearance is present on the database, the simulator is unaware of these products' existence to be sold. Moreover, by this date, allocation

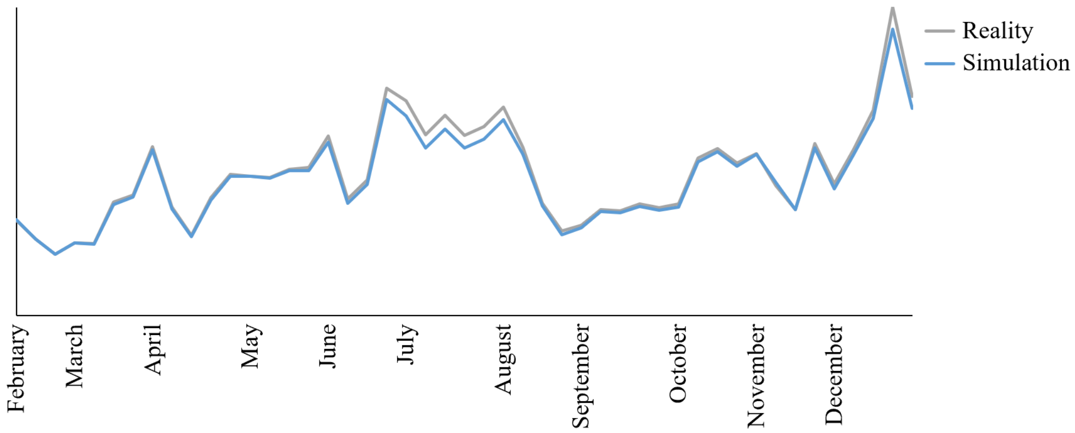

Figure 4.3: Comparison of sales quantities per week between simulation results and reality.

decisions that were not completely guided by the maxDemand variable add up, and inventory could be stuck in stores for which selling periods have ended. This is allowed by calculating the remaining demand for the time of the request and not for the time of the arrival of the product, which would completely fix shipments to perfect forecast conditions and is, therefore, avoided. It is important to mention that due to the lower prices practiced during these markdown periods, the global profit does not take as big a toll, and a plot for comparison would show two close to overlapping curves.

The second aspect to evaluate is concerned with the shipments between DC and stores. Overall results showed that 5.8% were not shipped on the simulation, which issued the need for further scrutiny. Figure 4.4 shows the comparison between the total number of pallets shipped per week on the simulation and on real data.

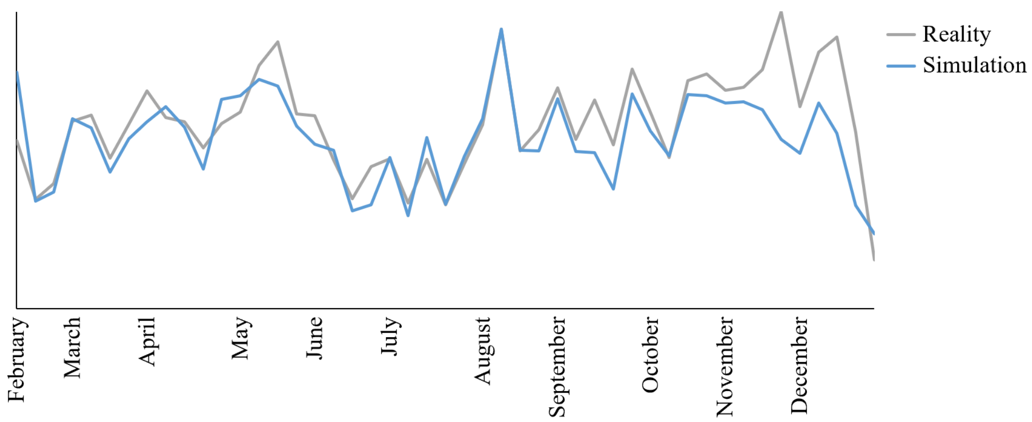

Figure 4.4: Comparison of number of pallets shipped per week between simulation results and reality.

For the majority of the period considered, the simulator does a reasonable job of mimicking reality, as the curves follow the same patterns. However, as the year gets close to the end, the gap widens, and a considerable amount of products seem to be missing from the deliveries. Two different factors were believed to be behind it. Firstly, since available DC inventory is constantly one month ahead, the last month could suffer from too tight restrictions on availability. This argument does not seem too strong when considering that sales do not seem to suffer much during this period. That being said, the fitting of inventory made possible by maxDemand could be to blame for this deviation, as no data for the year 2019 is considered, leading stores to end the year without considering keeping inventory to fulfill upcoming demand. As an extension for this model, final inventory levels could be implemented to decrease this terminating effect.

Besides the shortcomings provided, and after suitable causality was found, the simulation model is capable of establishing microscopic interactions that still make sense at an aggregate level. One important factor to bear in mind is that after credibility is established for the results, all comparisons can be done between simulation outputs, eliminating the effects aforementioned. Consequently, the model can be labeled as validated as all major relationships are correctly represented and experimentation can take place.

4.4 Metamodel

Although simulation provides extremely detailed outputs for each configuration tested, it comes with the drawback of being a computationally expensive method for experimentation of an extensive number of possibilities. In order to obtain interesting solutions in a timely manner, a metamodel was introduced which provides an estimation for the sales made possible by each delivery pattern, that is, how efficient a certain combination of shipment days is at reducing stock-outs for that specific store.

4.4.1 Model

Numerous methods are presented in the literature for this end but, undoubtedly, the most elementary one is by establishing a relationship between variables using a simple linear regression, so it should be evaluated if it is applicable in this case. Choosing to tackle a multitude of small problems instead of a single large one, each store was considered independent from the others, and a function was approximated per store. Although this is not entirely true due to limited centralized inventory (for instance, if a store only receives on Fridays it could always end up missing items that were already shipped to other stores), the simulation runs were defined in ways that balance equality through the stores, diminishing these effects which are assumed to be irrelevant.

The first step in the methodology is to choose a continuous variable that can effectively substitute discrete days enumeration on the classification of each scenario, that is, replace a Monday-Wednesday combination with a number: the weighted replenishment speed. Using an average week of deliveries for a scenario in which stores can receive inventory every weekday, one can estimate the effect that different delivery windows have, by delaying deliveries which are not possible on the current combination to the next shipment day. To turn this idea more comprehensible consider a Portuguese store that receives inventory every day except on Thursdays. This will result in the quantities scheduled for this day being pushed until Friday, as when OUTLs are checked with 48h lead time on Wednesday, they will already incorporate two days of sales since the last order was placed. The reasoning behind this logic becomes a bit trickier when talking about weekend replenishments, as there is the need to separate the quantities to correctly assess the days until inventory returned to a store after the purchase. That said, using the information in Table 3.1, one can estimate which percentage of the delivery quantity is related to each day, arriving at average quantities sold and replenished per day. By multiplying these quantities with each day's time until delivery, a global value is found for weighted replenishment speed for that combination of delivery days, for that store. It is important to mention that using shipment data instead of sales is extremely relevant due to the fact that not all items are replenished in a fast-fashion context, where collections are always changing. The described heuristic was used to find this characteristic for the two periods of the year under analysis.

After this independent variable is defined, simulated data was used for five different scenarios, with all stores possessing the same delivery frequency that ranged from two times to all days of the week (always not including the weekend) and for the current varied delivery patterns. These five data points were used to fit the model and find the relationship between replenishment speed and profit (meaning, the difference between the selling price and buying cost), for each store. The calculation of the parameters for the linear regression was done with Microsoft Excel. As an example, Figure 4.5 illustrates the data points obtained and the estimated trend for one store.

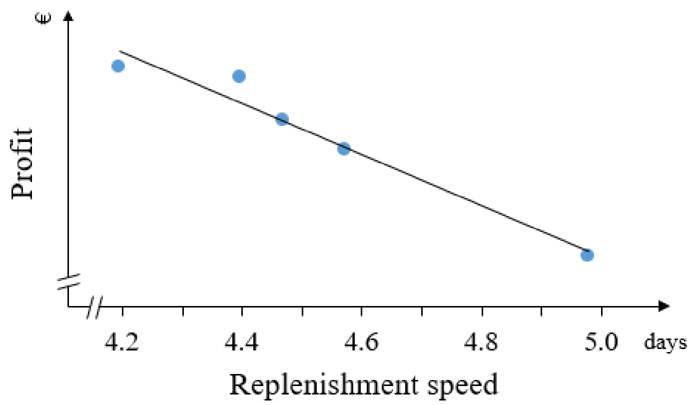

Figure 4.5: Example of the relationship between replenishment speed and obtained profit for one store.

4.4.2 Validation

In order to prove the linear regression's capability of providing a good substitution for the simulation model, a validation step must be performed (Wang and Shan, 2007). Predictive models can be tested for accuracy using metrics that compare the forecast values with real data, being the two most used ones (due to an easier evaluation of results) the Mean Average Percentage Error (MAPE) and Mean Percentage Error (MPE or Bias) (Schuller, 2018).

The most straightforward and first approach when using linear regression models is to determine the R², which averaged 86% for all stores and could be considered reasonable. MAPE averaged 0.07% for all data points, which seems a pretty small error, and Bias calculations were approximately 0%.

Since limited data is available, no extensive test and train methodologies can be used to reach better precision. However, the leave-one-out approach was used to provide insights into the validity of the model. Since common delivery frequency data points provide a strong base for the model by incorporating extreme values, the As-Is configuration was chosen to be only used as testing data. R² increased by 2% (which was expected as it is a measure that strongly depends on the number of points), and MAPE stayed below 0.1% for both test and training data.

It is clear that the low variation found between each scenario and the high absolute values are behind this apparently strong model. However, this makes this simplified approach suitable for the problem at hand and optimization can lean on these results to provide an approximation that will prioritize higher selling stores.

4.5 Optimization

The optimization module concludes the definition of improved delivery patterns, relying on the previous analysis to provide needed inputs. The main goal is to choose the best shipment days for each store while maintaining holistic integration. For the proposed methodology, we assume fixed total number of units shipped per week per store which, when considering the costs introduced in Table 3.2, means that only the ones varying with pallet number are affected by shipping policies. Keeping in mind that aggregating more days can potentially reduce the total number of pallets shipped per week for that store, a higher frequency is expected to yield higher logistic costs and higher sales (as it becomes more responsive to replenishment needs). As already mentioned, we chose not to include store operations, since employees are believed to have enough idle time to process receptions, meaning that changes do not directly impact expenses. The notation used is summarized in Table 4.3.

Firstly, all different scenarios per store $s \in S$ were listed, enabling the use of OpenSolver®, an Excel add-in for optimization. For each delivery pattern $n \in N$ and each weekday $d \in D$ (from Monday until Friday), a binary variable y_{dn} defined whether a shipment was allowed. Using an average week for a scenario with the highest delivery frequency, the approach introduced in the previous section was replicated for each store-delivery days combination (that is, each row), finding shipment quantities per day, ${q^R}_{sdn}$. This allowed each row to be evaluated for expected direct product profits P_{sn}, using the metamodel introduced in the previous section, transportation costs T_{sn} and expedition costs E_{sn}, by transforming number of products into number of pallets. The general equation for the calculation of P_{sn} using the already-defined linear regression is presented in equation 4.5, which takes as input the replenishment speed for each scenario and the α_s and β_s

parameters estimated by the metamodel.

$$P_{sn} = \alpha_s + \beta_s \cdot RVel_{sn} \qquad s \in S, n \in N \tag{4.5}$$

The incorporation of first shipments must not be forgotten, as they provide extra flexibility for the solution though with extra complexity to the model. Hence, two sets of decision variables are found: x_{sn}, which is a binary variable that states whether a delivery pattern was chosen for that a store, and NI_{sd}, as the percentage of new introductions to be shipped each weekday for each store. Let q_s^{NI} be the average quantity of new products sent every week to each store, λ the average number of units that fit into a box and ψ the number of boxes in a pallet, it is possible to define the number of pallets shipped each day for each store for every delivery pattern, Q_{sdn}, as:

$$Q_{sdn} = \left\lceil \frac{q_s^{NI} \cdot NI_{sd} + q_{sdn}^{R}}{\lambda \cdot \psi} \right\rceil \qquad s \in S, d \in D, n \in N \tag{4.6}$$

A ceiling needs to be implemented as even one box will be charged as a pallet in transportation and incur similar costs in expedition. Equations 4.7 and 4.8 provide the explicit calculation, being e the expedition costs per pallet and t_s the transportation costs per pallet for store s.

$$T_{sn} = \sum_{d \in D} Q_{sdn} \cdot t_s \quad s \in S, n \in N \tag{4.7}$$

$$E_{sn} = \sum_{d \in D} Q_{sdn} \cdot e \quad s \in S, n \in N \tag{4.8}$$

Altogether, the objective function can be described by equation 4.9, where we aim to maximize overall profits, accounting for costs of shipping, adding the individual estimation for each store, for the chosen delivery days.

$$\text{Maximize} \sum_{s \in S} \sum_{n \in N} (P_{sn} - T_{sn} - E_{sn}) \cdot x_{sn} \tag{4.9}$$

Certain constraints need to be added to the formulation to align the problem with reality. Firstly, capacity at the DC must be assured for each day CAP_d^{DC}, meaning that shipped quantities for all stores should be below the threshold, as described in 4.10. Moreover, due to space limitations at the stores, the number of pallets received must fit in the store receiving area CAP_s^{Rec} (equation 4.11).

$$\sum_{s \in S} \left(q_s^{NI} \cdot NI_{sd} + \sum_{n \in N} q_{sdn}^{R} \cdot x_{sn} \right) \leq CAP_d^{DC} \quad \forall d \in D \tag{4.10}$$

$$\sum_{n \in N} Q_{sdn} \cdot x_{sn} \leq CAP_s^{Rec} \quad \forall s \in S, d \in D \tag{4.11}$$

In order to align first introductions with business strategies, each store's percentage of clothing (apparel) items on new introductions App_s (that is, excluding accessories and shoes, for instance) must be guaranteed to reach the stores within the first two days of the week. This business restriction was included as defined in 4.12.

$$\sum_{d \in \{Monday, Tuesday\}} NI_{sd} \geq App_s \qquad \forall s \in S \tag{4.12}$$

At last, to keep decision variables within expected bounds, constraints 4.13 through 4.17 were included in the formulation, starting by limiting new introductions to days were the chosen policy has a delivery, forcing 100% of new introductions to be allocated per store and only defining one delivery pattern for each store. Additionally, NI_{sd} was made non-negative and x_{sn} binary.

$$NI_{sd} \leq \sum_{n \in N} y_{dn} \cdot x_{sn} \quad \forall s \in S, d \in D \tag{4.13}$$

$$\sum_{d \in D} NI_{sd} = 1 \quad \forall s \in S \tag{4.14}$$

$$\sum_{n \in N} x_{sn} = 1 \quad \forall s \in S \tag{4.15}$$

$$NI_{sd} \geq 0 \quad \forall s \in S, d \in D \tag{4.16}$$

$$x_{sn} \in \{0, 1\} \quad \forall s \in S, n \in N \tag{4.17}$$

Table 4.3: Table of notation for optimization.

Sets and indices	
$s \in S$	Stores
$d \in D$	Weekdays = {Monday, Tuesday, Wednesday, Thursday, Friday}
$n \in N$	Delivery patterns
Parameters	
t_s	Transportation costs per pallet for store s
e	Expedition costs per pallet
q^R_{sdn}	Replenishment quantities for store s on day d with delivery pattern n
q^{NI}_s	New introduction quantities for store s
CAP^{DC}_d	DC processing capacity (in items) on day d
CAP^{Rec}_s	Receiving capacity (in pallets) for store s
App_s	Apparel percentage on new introductions for store s
y_{dn}	Binary parameter that indicates whether delivery pattern n has delivery on day d
λ	Number of units per box
ψ	Number of boxes per pallet
α_s	Intercept of the sales profit regression for store s
β_s	Slope of the sales profit regression for store s
RVel_{sn}	Replenishment speed for store s and delivery pattern n
P_{sn}	Sales profit for store s and delivery pattern n
Auxiliary variables	
Q_{sdn}	Number of pallets delivered to store s on day d for delivery pattern n
T_{sn}	Transportation costs per week for store s and delivery pattern n
E_{sn}	Expedition costs per week for store s and delivery pattern n
Decision variables	
x_{sn}	Binary variable that indicates whether delivery pattern n is chosen for store s
NI_{sd}	Percentage of New Introduction quantities to be shipped on day d for store s

Chapter 5

Case Study Results

This chapter presents the results of applying the described methodology to the fashion retailer under analysis. Since no implementation was performed by the end of the study, a comparison basis will be used. Through simulation of one year with historical data, the policy currently in place is replicated and a baseline As-Is scenario is established.

A detailed solution was found for the optimization of delivery patterns, compiled in Section 5.1. Several validation approaches were used to evaluate scenarios. Additionally, in Section 5.2, the accomplished simulation model was used to experiment with an emergency shipment policy.

5.1 Delivery Pattern optimization

In order to provide a better understanding of the effects resulting from various delivery frequencies, Figure 5.1 illustrates the differences between costs in the 3 dimensions considered: DC, transportation and lost profits.

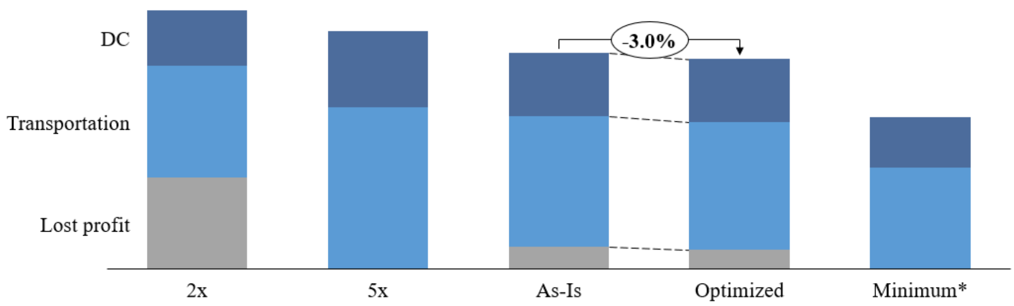

Figure 5.1: Comparison of the costs obtained for multiple scenarios tested.

In the '5x' scenario, all stores were supplied five times per week, corresponding to the highest delivery frequency possible, which was used as a threshold for the sales the current logistics and procurement structure is capable of capturing. An extremely high transportation expenditure is observed, in contrast with the one enabled by the '2x' option, where only two delivery windows were allowed per store and shipments become more efficient. However, the slower reaction to

Table 5.1: Comparison of the distribution of delivery frequencies for the As-Is and the Optimized scenario.

Delivery frequency	As-Is	Optimized	
		1st partition	2nd partition
2	21%	26%	4%
3	45%	50%	46%
4	19%	14%	37%
5	15%	10%	13%
Average	3.3	3.1	3.6

purchases and resulting amounts of lost sales (and consequently, lost profit) clearly outweigh the benefits. This result immediately leads to the conclusion that much lower delivery frequencies than the ones currently implemented will not be suited for the specific circumstances.

As a way to provide a theoretical target for the best case, the 'Minimum*' costs were defined by assuming the possibility to fulfill all the sales of the '5x' scenario by a single delivery to each store per week. It is clear that this might not feasible due to store or DC capacity constraints (and would represent a better ability to forecast demand), yet it provides the awareness that even in the best case scenario there is a considerable fixed cost structure required.

Following the methodology presented in the previous chapter, an 'Optimized' set of delivery patterns was found, splitting the year into two partitions divided at the beginning of November. When comparing global costs to the As-Is configuration, a 3.0% decrease was found (with reductions in all dimensions considered), in addition to a 0.1% reduction in inventory left unsold at the end of the simulation. The fact that sales were increased and shipment costs lowered indicates that better allocation of delivery windows for each store was obtained. Similarly to the 3.3 average delivery frequency currently used, the solution reached 3.1 for the first part of the year and 3.6 for the peak period afterwards. This goes in line with expectations that higher availability during this time might be worth the extra costs incurred. The shift to increasing trips to stores is noticeable in Table 5.1, as very few stores get deliveries only two times per week at the end of the year (only 4%). It is interesting to notice how As-Is values are always higher for the maximum delivery frequency, which by itself could steer company stakeholders towards a smaller number of shipments.

It is important to guarantee that capacity requirements for preparing shipments at the DC and at the receiving end, in the store, are within acceptable bounds. Since additional workforce can be scheduled for peak periods (such as new store openings or Christmas shopping sprees) and stores can occasionally create room when really needed, the As-Is situation was used as a proxy for the acceptable amount of irregular situations. Figure 5.2 presents the percentage of products that were over theoretical capacity at the store, CAP_s^{Rec}, and at the DC, CAP_d^{DC}.

For the first dimension, all scenarios perform similarly low, with the exception of the lowest delivery frequency, '2x', which could potentially cause problems due to repeated pressure on storage space. The 'Optimized' alternative stayed close to the results for the 'As-Is' configuration, proving

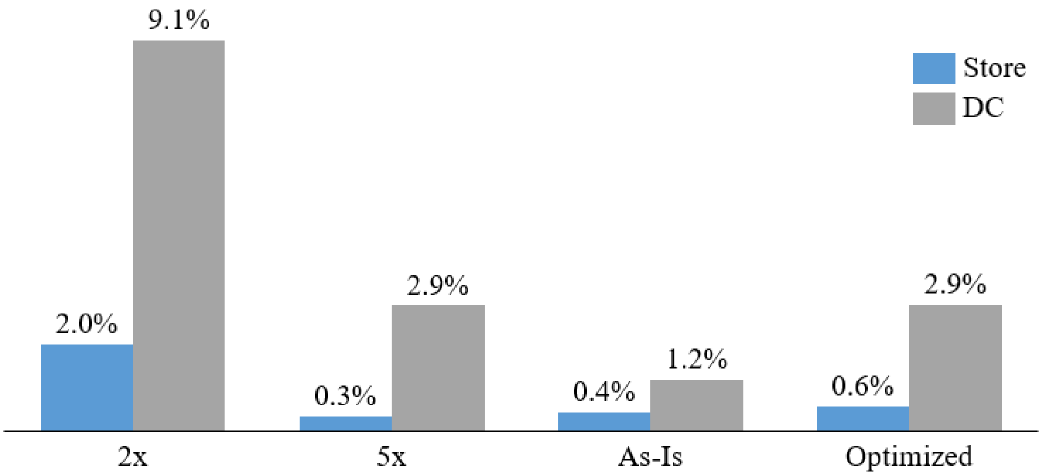

Figure 5.2: Comparison of the percentage of units over capacity at store reception and DC preparation for multiple scenarios tested.

itself as acceptable. However, a more strenuous landscape is shown at the DC, where the sorter capacity is not capable of fulfilling shipments of 9,1% of the units for the ’2x’, when considering current timetables, dismissing this scenario completely. The results for the comparison between the other three lead to an interesting conclusion: since store delivery patterns were aimed at improving replenishment velocity, most stores got delivery windows on ’important’ days, that is, the quicker possibility of restocking weekend purchases (which makes the ’5x’ and the ’Optimized’ solutions so similar). Although this did not prove problematic for the average week considered in the optimization module, an increase in the percentage of units over-capacity is evidenced for the full year run. It is not straightforward that this would block this scenario, as alternative solutions could be found such as increasing sorter hours, better allocation of first shipments, manual sorting, among others. Nevertheless, the extra entropy should weight on the decision to implement the changes.

Taking into consideration the mentioned limitations, the cost reduction obtained for the ’Optimized’ scenario was, however, not considered enough to justify the necessary adjustments by company stakeholders, due to the proximity to current practices. A positive note still originated from the experimentation, as confidence in the current empirical method was increased and the need for implementing changes demystified.

5.2 Emergency shipments simulation

As presented in Section 4.3, an important feature was added to the simulator, enabling experimentation with a shorter lead time. Since current practices hold inventory in the DC for around one day and a half between store ordering and shipment (which takes place at the end of the afternoon), the impact of being more agile in this preparation, and shipping within the day of ordering, showed potential for assessment. Following, a further explanation of the reasoning behind the approach and the results obtained from simulation are presented.

5.2.1 Contextualization

The DC starts processing store orders in the morning of the requests, and after consolidated picking, the sorter begins separating products per store. However, due to the large volume to be processed, this process ends after carrier's departure time, meaning shipment will only take place on the following day (refer back to Section 3.1 for more details). In order to allow ending preparation in time, a smaller subset of products can be chosen to be prioritized and rushed through operations, which, in line with the availability mindset, were defined as stock-out products at each store.

The current replenishment policy is merely responsive to the OUTL, which is checked for every product whenever there is a 'request' window, in line with the lead time required to comply with the delivery days defined. However, for the testing hypotheses at hand, and to increment possibilities of sales, all items were checked every day for zero inventory. This results in shipments happening most days for every store, increasing transportation costs. A new carrier was considered for these trips, which presented tariffs based on box quantities instead of pallets.

This operational change came at the expense of some activities previously done at the DC (and for that reason, until this point, out of scope for the analysis), which now need to be pushed to the store, such as tagging individual products. This important adjustment means that costs at this dimension should be evaluated for holistic calculation, as well as some extra entropy at the DC from reducing economies of scale.

The main goal of the analysis is then to understand whether the extra sales made possible by quicker replenishment velocities are enough to outweigh the bigger operational costs while keeping current delivery patterns.

5.2.2 Results

The final results for the simulation run allowing Emergency Shipments (ES) are depicted in Figure 5.3, in comparison to the As-Is scenario. As mentioned, new dimensions were incorporated in the analysis to fully evaluate the consequences of the alternative.

Reducing replenishment lead time of critical SKUs did, as expected, increase sales, creating a new baseline for possible purchases and putting in perspective previously obtained values for lost profits. A closer examination of the decomposition of the bars reveals how costs shift from the DC dimension to store operations, due to customization activities happening partially at this stage. The consequent reduction in DC expenses is, however, neutralized by the expected increase in shipments preparation, whose number is now substantially higher. Transportation also takes a toll due to this fact. Nevertheless, it was possible to determine that additional captured demand is indeed sufficient to outweigh these extra logistic expenses, leading to a 6.9% overall decrease in the considered costs for a year of simulation. Moreover, 1.1% less inventory was remaining after the simulation run.

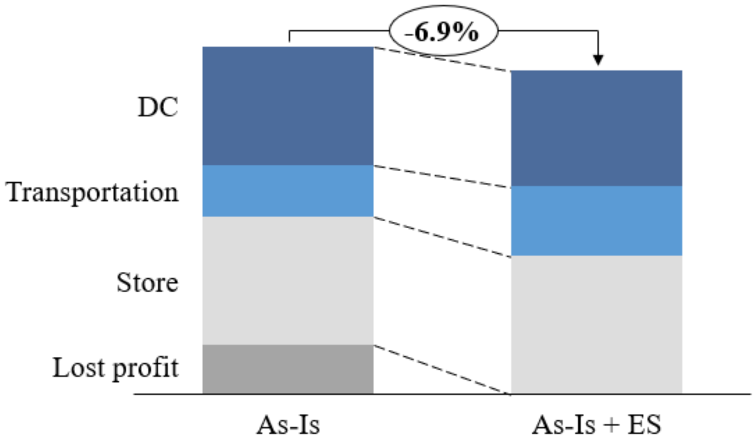

Figure 5.3: Comparison of the costs obtained with or without emergency shipments.

In order to validate the results obtained, received quantities (at the store) per weekday were plotted against current ones, differentiating between emergency (in white) and regular (in blue) requests, which is presented in Figure 5.4.

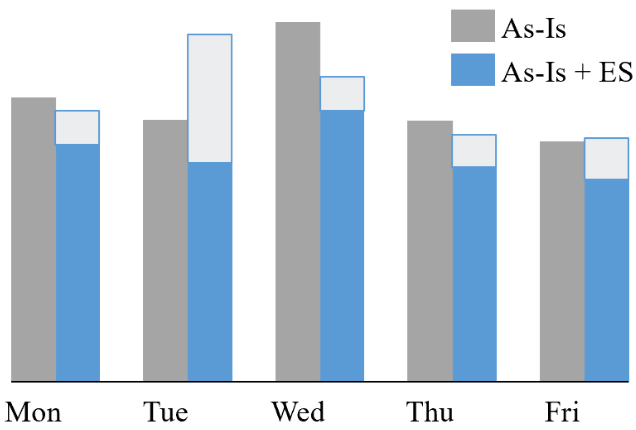

Figure 5.4: Comparison of the average quantities received for each day of the week with or without emergency shipments.

As would be expected, a large proportion of these priority products is received on Tuesdays, resulting from stock-outs reported on Monday, after weekend sales. This reduces the number of products arriving in stores Wednesday since some are sped up through DC processes. Overall, ES account for 19% of units delivered, which issues further investigation into the feasibility of the scenario. Out of the 332 days simulated, 14 were marked as over-capacity for the special sorter configuration needed to include emergency shipments. To evaluate the impact of this constraint, a sensitivity analysis is pictured in Figure 5.5 with regard to hour variations in operation time.

An initial remark should be done on how the '+0' available time was defined in a conservative manner, ensuring possibility of execution. Consequently, reducing one hour would represent an

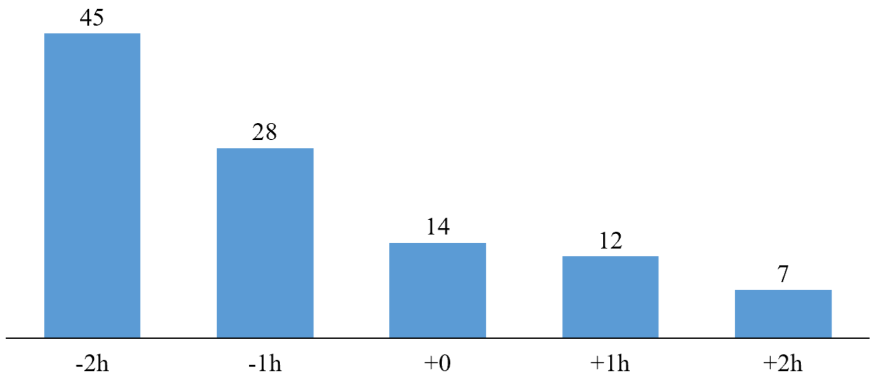

Figure 5.5: Sensitivity analysis of the impact of changing sorter operating time for the emergency shipments scenario.

extremely unlikely scenario, although with the duplication of complicated days. However, if two additional hours are allocated to sorting products, the value drops to half.

As for the store dimension, no additional strain is incurred in relation to the As-Is configuration as shipments are performed in boxes for days without scheduled deliveries and per-store quantities are small and mostly do not impose changes in pallet number for the remaining days.

The relative simplicity of implementation of this policy attracted company decision-makers to the pursuit of this path, which should be ongoing in the near future.

Chapter 6

Conclusions and Future Research

This dissertation focused on establishing store delivery patterns within the context of a fashion retailer in an integrative approach. To do so, dependencies between supply chain dimensions were analyzed, reaching a verdict on relevant restrictions and costs for the distribution center, transportation and store subsystems. By mapping these impacts, a holistic analysis is obtained and recommendations can be confidently introduced.

The proposed methodology comes as a consequence of the inherent characteristics of the apparel industry, where short product life-cycles lead to a constant change in assortment, with new items being sent to stores weekly. Moreover, the availability-driven demand weights on the need for assuring a quick replenishment of shelves, in order to avoid stock-outs. A product not sold in a customer trip can end up in markdowns or be left unsold at the end of the season, with a reduction of profits. Moreover, very low predictability requires a fast reaction to sales, as forecasts come with a high error margin at the SKU-store level.

Although complex methods have been proposed for the optimization of grocery delivery days (which include modelling effects of order sizes on various activities), none have yet tackled the influence these might have on stock-outs. This represents the main addition to the research field, and the negative consequence to balance with lower expedition costs. To do so, a microscopic approach becomes necessary as, otherwise, the relationship between lost profits and decision variables could not be assessed. After estimating demand for days where it was not recorded, a simulator is introduced, detailing the impacts of distinct replenishment policies on store sales. However, it would be an extremely time-consuming process to aim for experimenting with all the possible combinations of store-delivery patterns, so a shortcut is brought on by incorporating a metamodel. With this tool, the optimization module could present reasonable results and arrive at ideal solutions for an average week. Validating the scenario then required the simulation run of the policy.

For the case study analyzed, a 3% cost improvement to current practices was found, though at the expense of slightly tighter logistic operations. Despite a substantial gain not being clear, company stakeholders believed the study still provided important insights for evaluating established systems. In order to leverage the development of the simulation model, an emergency shipments policy, which added to current delivery patterns, was tested and results were remarkably promising for implementation possibility, with 6.9% of savings in costs incurred.

The simulation-optimization sequence suggested should not be interpreted as a solution for a specific retailer problem but rather as a generic tool with the potential to answer a complex analytic challenge of any supply chain facing identical interactions. Significant improvement opportunities could be estimated for a less established player, which has not yet been able to tweak operations into an equilibrium, since the methodology did, in fact, propose a solution close to minimum costs.

Several limitations on the assumptions made were mentioned throughout the dissertation and could instigate further research. For one, the approach used for steering of inventory could be substituted by a more reasonable one, such as fitting a last delivery date. It is clear, however, that being able to estimate demand out of historical periods could be enough to eliminate the correction completely. Additionally, the simplistic metamodel incorporated could be made more reliable by more sophisticated prediction techniques. Moreover, for the customer-centric approach considered, store costs were mainly disregarded in optimization, which could prove not to be the path for all cases. The transportation was also straightforward due to outsourcing to a 3PL. Additional research could combine a routing problem to the methodology, reaching efficient costs for a company-owned fleet of trucks.

The power of simulation for impact assessment was only slightly harnessed and additional features could be added for evaluation of alternatives. For instance, lateral transshipments between stores could represent a more efficient way of performing emergency deliveries. Furthermore, OUTL calculations could be redefined and even fully questioned, as they have been defended as far from optimal for lost sales models (Bijvank et al., 2015). Allocation decisions could also be incorporated in the simulator, in contrast to the pure first-come, first-served approach used.

Finally, as a closing remark, it is important to stress the relevance of the research on the ever-growing fashion industry. It should not be seen as a pure consequence of success that Zara outscores any other company on the number of published articles, but rather as a contributing factor to it. Pertinent studies on this unexplored field can set the difference between winning or losing a purchase, a client or the market itself.

References

Aftab, M. A., Yuanjian, Q., Kabir, N., and Barua, Z. (2018). Super Responsive Supply Chain: The Case of Spanish Fast Fashion Retailer Inditex-Zara. *International Journal of Business and Management*, 13(5):212–227.

Al-Zubaidi, H. and Tyler, D. (2004). A simulation model of quick response replenishment of seasonal clothing. *International Journal of Retail & Distribution Management*, 32(6):320–327.

Amed, I., Balchandani, A., Beltrami, M., Berg, A., Hedrich, S., and Rölkins, F. (2018). The State of Fashion 2019. Technical report.

April, Glover, Kelly, and Laguna (2003). Practical introduction to simulation optimization. In *Proceedings of the 2003 Winter Simulation Conference, 2003.*, volume 1, pages 71–78 Vol.1.

Banks, J., Carson II, J. S., Nelson, B. L., and Nicol, D. M. (2010). *Discrete-Event System Simulation*. Pearson, New Jersey, 5th edition.

Bartmann, D. and Beckmann, M. J. (1992). *Inventory Control: Models and methods*. Springer-Verlag, Berlin Heidelberg, 1st edition.

Bijvank, M., Bhulai, S., and Tim Huh, W. (2015). Parametric replenishment policies for inventory systems with lost sales and fixed order cost. *European Journal of Operational Research*, 241(2):381–390.

Cardós, M. and García-Sabater, J. P. (2006). Designing a consumer products retail chain inventory replenishment policy with the consideration of transportation costs. *International Journal of Production Economics*, 104(2):525–535.

Caro, F. and Gallien, J. (2010). Inventory Management of a Fast-Fashion Retail Network. *Operations Research*, 58(2):257–273.

Caro, F. and Gallien, J. (2012). Clearance Pricing Optimization for a Fast-Fashion Retailer. *Operations Research*, 60(6):1404–1422.

Carson, Y. and Maria, A. (1997). Simulation Optimization: Methods And Applications. In *Winter Simulation Conference Proceedings,*, pages 118–126.

Cerda, C. B. R. and Monteros, A. J. E. (1997). Evaluation Of A (R,s,Q,c) Multi-item Inventory Replenishment Policy Through Simulation. In *Winter Simulation Conference Proceedings,*, pages 825–831.

Chartniyom, S., Lee, M. K., Luong, L., and Marian, R. (2007). Multi-Location Inventory System with Lateral Transshipments and Emergency Orders. In *2007 IEEE International Conference on Industrial Engineering and Engineering Management*, pages 1594–1598.

Christopher, M., Lowson, R., and Peck, H. (2004). Creating agile supply chains in the fashion industry. *International Journal of Retail & Distribution Management*, 32(8):367–376.

Christopher, M. and Peck, H. (1997). Managing Logistics in Fashion Markets. *The International Journal of Logistics Management*, 8(2):63–74.

Collado, P. B. and Martínez-de Albéniz, V. (2017). Estimating and Optimizing the Impact of Inventory on Consumer Choices in a Fashion Retail Setting. pages 1–28.

Correa, J. R. J. R. (2007). *Optimization of a fast-response distribution network*. PhD thesis, Massachusetts Institute of Technology.

Corsten, D. and Gruen, T. (2003). Desperately seeking shelf availability: An examination of the extent, the causes, and the efforts to address retail out-of-stocks. *International Journal of Retail & Distribution Management*, 31(12):605–617.

De Marco, A., Cagliano, A. C., Nervo, M. L., and Rafele, C. (2012). Using System Dynamics to assess the impact of RFID technology on retail operations. *International Journal of Production Economics*, 135(1):333–344.

Duong, L. N. K. and Wood, L. C. (2018). Discrete Event Simulation in Inventory Management. In Mehdi Khosrow-Pour, D. B. A., editor, *Encyclopedia of Information Science and Technology, Fourth Edition*, pages 5335–5344. IGI Global, Hershey, PA, USA.

Fernie, J. and Grant, D. B. (2014). Investigating on-shelf availability in the UK retail clothing sector. In Choi, T.-M., editor, *Fast Fashion Systems: Theories and Applications*, chapter 7, pages 95–110. CRC Press, Florida.

Figueira, G. and Almada-Lobo, B. (2014). Hybrid simulation-optimization methods: A taxonomy and discussion. *Simulation Modelling Practice and Theory*, 46:118–134.

Fraiman, N. M., Singh, M., Paris, C., and Arrington, L. (2010). *Zara*. Columbia CaseWorks, Columbia Business School.

Gallien, J., Mersereau, A., Garro, A., Mora, A., and Vidal, M. (2015). Initial Shipment Decisions for New Products at Zara. *Operations Research*, 63(2):269–286.

Garro, A. (2011). *New product demand forecasting and distribution optimization: a case study at Zara*. PhD thesis, Massachusetts Institute of Technology.

Gaur, V. and Fisher, M. L. (2004). A Periodic Inventory Routing Problem at a Supermarket Chain. *Operations Research*, 52(6):813–822.

Hachicha, W., Afli, A., and Elleuch, F. (2013). Evaluating emergency lateral transshipment policies using simulation-based approaches. In *2013 International Conference on Advanced Logistics and Transport*, pages 470–475.

Holzapfel, A., Hübner, A., Kuhn, H., and Sternbeck, M. G. (2016). Delivery pattern and transportation planning in grocery retailing. *European Journal of Operational Research*, 252(1):54–68.

Huang, H. and Liu, Q. (2017). Intelligent retail forecasting system for new clothing products considering stock-out. *Fibres and Textiles in Eastern Europe*, 1(121):10–16.

Iannone, R., Ingenito, A., Martino, G., Miranda, S., Pepe, C., and Riemma, S. (2013). Merchandise and Replenishment Planning Optimisation for Fashion Retail. *International Journal of Engineering Business Management*, 5:26.

Jalal, H., Dowd, B., Sainfort, F., and Kuntz, K. M. (2013). Linear regression metamodeling as a tool to Summarize and present simulation model results. *Medical Decision Making*, 33(7):880–890.

Jalali, H. and Nieuwenhuyse, I. V. (2015). Simulation optimization in inventory replenishment: a classification. *IIE Transactions*, 47(11):1217–1235.

Jayaraman, R. and Tabucanon, M. T. (1984). Co-ordinated Versus Independent Replenishment Inventory Control. *International Journal of Operations & Production Management*, 4(1):61–68.

Kleijnen, J. P. (2005). An overview of the design and analysis of simulation experiments for sensitivity analysis. *European Journal of Operational Research*, 164(2):287–300.

Kotzab, H. and Teller, C. (2005). Development and empirical test of a grocery retail instore logistics model. *British Food Journal*, 107(8):594–605.

Kuhn, H. and Sternbeck, M. G. (2013). Integrative retail logistics: An exploratory study. *Operations Management Research*, 6(1-2):2–18.

Liao, Y., Shen, W., Hu, X., and Yang, S. (2014). Optimal responses to stockouts: Lateral transshipment versus emergency order policies. *Omega (United Kingdom)*, 49:79–92.

Longo, F. and Mirabelli, G. (2008). An advanced supply chain management tool based on modeling and simulation. *Computers and Industrial Engineering*, 54(3):570–588.

Martino, G. (2015). *Supply Chain Management in the Fashion Retail Industry: a multi-method approach for the optimisation of performances*. PhD thesis, University of Salerno.

Martins, S. S. B. (2018). *Retail Distribution Planning to Brick-and-Mortar Stores*. PhD thesis, Faculdade de Engenharia da Universidade do Porto.

Mattila, H., King, R., and Ojala, N. (2002). Retail performance measures for seasonal fashion. *Journal of Fashion Marketing and Management: An International Journal*, 6(4):340–351.

Mou, S., Robb, D. J., and DeHoratius, N. (2018). Retail store operations: Literature review and research directions. *European Journal of Operational Research*, 265(2):399–422.

Musalem, A., Olivares, M., Bradlow, E. T., Terwiesch, C., and Corsten, D. (2010). Structural Estimation of the Effect of Out-of-Stocks. *Management Science*, 56(7):1180–1197.

Myers, B. R. (2009). *Analysis of an inventory system with product perishability and substitution: A simulation-optimization approach*. PhD thesis, Drexel University.

Nieves-Rodríguez, E. B., Cao-Alvira, J. J., and Pérez, M. M. (2015). The Influence of Special Occasions on the Retail Sales of Women's Apparel. In Kubacki, K., editor, *Proceedings of the Academy of Marketing Science*, pages 213–221, Cham. Springer International Publishing.

Nueno, J. L. and Ghemawat, P. (2003). ZARA: Fast Fashion. *Harvard Business School Cases*, 1:1–35.

Nuttle, H. L. W., King, R. E., Hunter, N. A., Wilson, J. R., and Fang, S. C. (2000). Simulation Modeling of the Textile Supply Chain. Part I: The Textile-plant Models. *The Journal of The Textile Institute*, 91(1):35–50.

Ovezmyradov, B. and Kurata, H. (2019). Effects of customer response to fashion product stockout on holding costs, order sizes, and profitability in omnichannel retailing. *International Transactions in Operational Research*, 26(1):200–222.

Pettinger, L. (2004). Brand culture and branded workers: Service work and aesthetic labour in fashion retail. *Consumption Markets & Culture*, 7(2):165–184.

Pires, M., Pratas, J., Liz, J., and Amorim, P. (2017). A framework for designing backroom areas in grocery stores. *International Journal of Retail & Distribution Management*, 45(3):230–252.

Reiner, G., Teller, C., and Kotzab, H. (2013). Analyzing the efficient execution of in-store logistics processes in grocery retailing - The case of dairy products. *Production and Operations Management*, 2(4):924–939.

Ronen, D. and Goodhart, C. A. (2008). Tactical store delivery planning. *Journal of the Operational Research Society*, 59(8):1047–1054.

Rouwenhorst, B., Reuter, B., Stockrahm, V., Van Houtum, G. J., Mantel, R. J., and Zijm, W. H. (2000). Warehouse design and control: Framework and literature review. *European Journal of Operational Research*, 122(3):515–533.

Schuller, P. (2018). *A machine learning approach to promotional sales forecasting*. Master's dissertation, Faculdade de Engenharia da Universidade do Porto.

Shannon, R. (1975). *Systems Simulation: The Art and Science*. Prentice-Hall, 1st edition.

Snyman, J. A. and Wilke, D. N. (2018). *Practical Mathematical Optimization: Basic Optimization Theory and Gradient-Based Algorithms*. Springer International Publishing, 2nd edition.

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(1):205–224.

Tagaras, G. and Vlachos, D. (2001). A Periodic Review Inventory System with Emergency Replenishments. *Management Science*, 47(3):415–429.

Ton, Z., Corsi, E., and Dessain, V. (2010). *Zara: Managing Stores for Fast Fashion*. Harvard Business School Publishing, Boston, MA.

Vulcano, G., van Ryzin, G., and Ratliff, R. (2012). Estimating Primary Demand for Substitutable Products from Sales Transaction Data. *Operations Research*, 60(2):313–334.

Wang, G. G. and Shan, S. (2007). Review of Metamodeling Techniques in Support of Engineering Design Optimization. *Journal of Mechanical Design*, 129(4):370–380.

Waters, D. (1992). *Inventory Control and Management*. Wiley, Chichester, 2nd edition.

Wecker, W. E. (1978). Predicting Demand from Sales Data in the Presence of Stockouts. *Management Science*, 24(10):1043–1054.

Yang, B. and Burns, N. (2003). Implications of postponement for the supply chain. *International Journal of Production Research*, 41(9):2075–2090.

Zinn, W. and Liu, P. C. (2008). A comparison of actual and intended consumer behavior in response to retail stockouts. *Journal of business logistics*, 29(2):141–159.