Optimizing Order-Routing Decisions: Leveraging Omni-Channel Supply Chain Fulfillment

by

Federico Markowicz

BSc. Mechanical Engineering
Technion, Israel Institute of Technology, 2011

Submitted to the MIT Sloan School of Management and Mechanical Engineering Department in partial fulfillment of the requirements for the degrees of

Master of Business Administration
and
Master of Science in Mechanical Engineering

In conjunction with the Leaders for Global Operations program at the
Massachusetts Institute of Technology
June 2017

Signature of Author: Signature redacted
Mechanical Engineering Department, MIT Sloan School of Management
May 12, 2017

Certified by: Signature redacted
Stanley Gershwin, Thesis supervisor
Senior Research Scientist at the MIT Department of Mechanical Engineering

Certified by: Signature redacted
Stephen Graves, Thesis supervisor
Professor of Management Science, MIT Sloan School of Management

Accepted by: Signature redacted
Rohan Abeyaratne
Chairman, Committee on Graduate Students
Department of Mechanical Engineering

Accepted by: Signature redacted
Maura Herson
Director, of MBA Program, MIT Sloan School of Management

This page intentionally left blank

Optimizing Order-Routing Decisions: Leveraging Omni-Channel Supply Chain Fulfillment

by

Federico Markowicz

Submitted to the MIT Sloan School of Management and the Mechanical Engineering Division on May 12, 2017 in partial Fulfillment of the Requirements for the Degrees of Master of Business Administration and Master of Science in Mechanical Engineering

Abstract

This thesis provides a deep mathematical analysis of the diverse alternatives for routing models considering an Omni-channel supply chain. The natural evolution of supply chains from traditional brick-and-mortar stores to an omni-channel supply chain, encompassing and merging e-commerce together with a multi-channel concept, allows businesses to reach new levels of operational efficiency by leveraging inventory closer to the customer and making decisions on the fly on how to better and more cheaply provide a service/product to the final consumer.

The flexibility and benefits, unfortunately, do not come without a certain dose of complexity and further development of the supply chain tactical implementation and systems. New alternatives to fulfill customer orders are available, which require greater screening among the different alternatives.

An effective routing model becomes essential to make sure these alternatives are properly considered in order to satisfy both the consumer and retailer objectives, such as on-time delivery of orders, retail stores' service levels, and fulfillment costs.

Thesis Supervisor: Stephen Graves
Title: Professor of Management Science, MIT Sloan School of Management

Thesis Supervisor: Stanley Gershwin
Title: Senior Research Scientist at the MIT Department of Mechanical Engineering

This page intentionally left blank

Acknowledgements

First, I would like to thank the entire Supply Chain Strategy and Analytics team at Verizon Wireless for the amazing support and guidance throughout my internship in Basking Ridge, especially to Yanai Golany and Vedat Akgun that provided the environment and tools for me to succeed.

In addition, I want to give a special mention to all the groups across Verizon Wireless Supply Chain organization for providing the expertise, guidance, and data to help me materialize my analysis. It was a pleasure interacting and working with everyone.

I'm very grateful that I had the opportunity to work and get to know such amazing people and professionals as: Viju Menon, Anne Robinson, Julian Sassoon, John Zuch, Benjamin Lowe, Jeremy Giese, Sagar Shah, Dave Allen, Alek Liskov, Jackie Kistler, and John Malayter.

I would like to acknowledge the LGO staff for their support and guidance during not only the internship, but also the entire LGO program.

I would also like to thank my thesis advisors, Prof. Steve Graves and Dr. Stan Gershwin. Both were a great source of knowledge and inspiration during my project and did a great job mentoring this work.

Finally, I would like to send my appreciation to Verizon Fios soccer team, 2016 Corporate tournament champions. It was enjoyable to join the team and friendly compete together with you, while having the opportunity to feel Verizon's culture outside the office desk.

Thank you all!

This page intentionally left blank

Table of Contents

List of Figures

List of Abbreviations

Abbreviation	*Explanation*
WC	Working Capital
SFS	Ship From Store
IT	Information Technology
SKU	Stock-keeping unit
DC	Distribution Center
DTC	Direct to Consumer
KPI	Key Performance Indicator
B2B	Business-to-Business
DOS	Days of Supply
SL	Service Level
FIFO	First in, first out
LP	Linear Program
ATP	Available to Promise

This page intentionally left blank

1 Introduction

Supply chain flexibility and efficiency are essential to reduce working capital (WC) and operational costs. The working capital represents the operating liquidity available to a business, and it is essential for companies to operate and conduct their business. In this paper when working capital is referred, it means inventory. Reducing inventory levels, or WC, will allow more liquidity to the company.

Verizon, like many retailers worldwide, are enhancing their omni-channel supply chain leveraging diverse distribution methods to reach customers. A network of large stores — with high overhead costs — can become a liability rather than an asset in recent years. Amazon, which has no stores, won market share with lower prices and huge selection. But retailers have begun fighting back by using technology to get more sales out of stores — and ship from store (SFS - i.e. the stores as potential locations to fulfill online orders) is a big part of the effort [1].

Having several hundreds of retail stores nationwide, Verizon holds a large amount of inventory in these locations that can be leveraged to obtain benefits for the company. Consequently, Verizon is evaluating shipping directly from its stores, adding flexibility to its supply chain. However, this broad flexibility also adds a large amount of complexity to the supply chain in aspects such as shipping logistics, stores' service levels, IT capabilities, order routing, etc.

Order routing, the focus of this thesis, becomes extremely complex when so many combinations of SKU-Store are being considered, in addition to other variables to be explained in the following sections. If we consider Verizon's current situation, where multiple distribution centers (DCs) are currently responsible to ship all the orders to customers, enabling SFS will elevate the routing puzzle to a much higher level of complexity, leaving a large space to optimize and achieve great results for the company.

1.1 Brief Statement of the Problem

Verizon Wireless reaches its customers for devices and accessories mainly through three channels: retail stores, direct to consumers (DTC), and indirect. The DTC orders arrive mainly via Verizon's website or telephone sales. Before SFS, the orders were shipped from one of the DCs directly to the customers, based on simple business rules. If SFS is fully implemented, multiple hundreds of stores will be able to fulfill customer orders. Deciding how to route the order, and which store/DC will fulfill it becomes a major challenge in several Key Performance Indicators (KPI): service levels, operational costs, WC, IT overload, etc.

Designing an optimal and effective routing model for this complex environment is the focus of this paper.

1.2 Verizon Communications & Verizon Wireless

Verizon Communications Inc. is a telecommunications company and the largest U.S. wireless provider. Its wide range of products extends from communications and information to entertainment. The three main businesses reported by Verizon are: Wireless, Wireline, and Enterprise.

- Wireless - provides mobile phone services for phones, tablets, and hotspot devices.
- Wireline – phone services, Internet access, and television to residences and small businesses.
- Enterprise – provides cloud-based solutions for corporate and government entities, delivering security, mobility, and information-sharing solutions.

Verizon employs more than 162,000 employees and its revenues for 2016 were $125.98 billion [2]. This project focuses on the Wireless business, which accounts for over 70% of the company's revenues.

Wireless provides mobile communications operating a nationwide 4G LTE network covering 98% of the U.S. population. As of October 2016, Verizon Wireless counted 144 million subscribers [3], maintaining its position as the leader in the U.S. market.

1.3 Verizon Global Supply Chain – Wireless Division

Verizon Global Supply Chain organization is responsible for ensuring the proper supply of products and services to satisfy customer expectations and satisfaction. More specifically, the Wireless Supply Chain is responsible for providing devices and accessories to enable services to its customers. By devices we mainly mean phones and tablets, and by accessories a wide range of products such as phone cases and protectors, Bluetooth headphones and speakers, batteries and chargers, etc.

There are three main channels in the Wireless division:

1. Retail – Verizon Wireless stores throughout the country.
2. Direct to Consumer (DTC) – mainly online and tele-sales.
3. Indirect – Through third party intermediaries (big wholesalers such as Walmart, Best Buy, etc.)

Experiencing and projecting an extensive growth in the DTC market, as customers more often order online, Verizon Wireless is investing in expanding and enhancing its supply chain with several projects. Each of them is aimed at improving customer experience and decreasing operational costs by allowing a more flexible and responsive supply chain.

Verizon's Global Supply Chain organization is aiming to enhance its omni-channel supply chain capabilities, where satisfying customer orders is not directly related to a specific channel, but allows full flexibility between these.

SFS is a wide and complex project, as it provides the opportunity to fulfill direct customers' orders from multiple locations, which implies not only supply-chain challenges, but also supporting operational processes at the stores as well as at the DCs.

1.4 Direct-to-Consumer Fulfillment

DTC as a channel represents a strategic focus for Verizon Supply Chain. With the consolidation of e-commerce in the last few years, the volume expected on this channel is expected to grow rapidly in the near future. DTC orders at Verizon are composed of several sub-channels, with online currently representing a significant portion of the total volume. Other sub-channels are: telesales, customer support, B2B, and store.

Before SFS was considered, the DTC fulfillment could be described by a simple logic that consisted of designating a primary and secondary DC per customer zone as responsible for the shipment. Section 4.2 details the process flow of current DTC fulfillment. An overview of the current state and the idea behind omni-channel is illustrated on Figure 1.

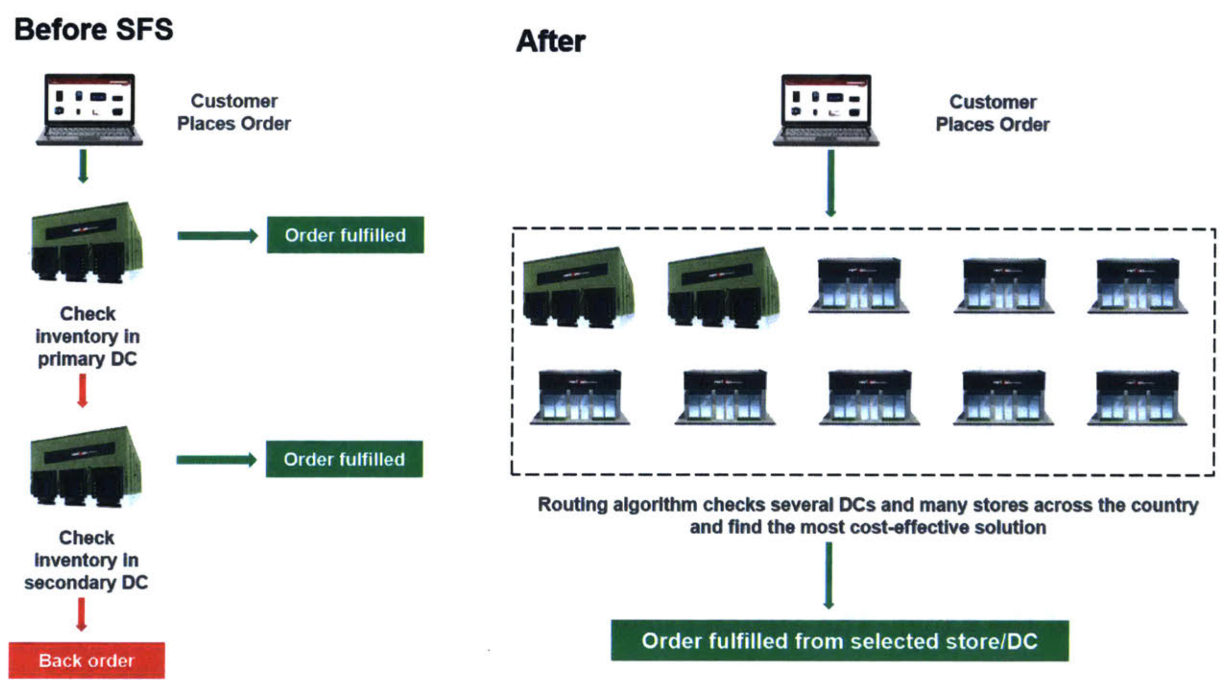

Figure 1 - DTC fulfillment overview

1.5 Thesis Overview

This thesis consists of the findings, insights and recommendations generated during a six-months internship with Verizon.

It starts by analyzing the current situation of Verizon Wireless' supply chain and how new ventures within the company disrupt and change it. It continues in Section 2 by providing an overview of the literature behind an omni-channel supply chain, analyzing industry trends and routing models, as well as the opportunities and challenges in implementing it.

Section 3 explains the methodology used to analyze the problem, guiding through the natural path followed by the company to gradually implement SFS with minimal disruption. A qualitative analysis of the current and future state is provided, together with an introduction to the main factors considered in the final analytical solution.

In Section 4 the routing models are presented in detail, including the rationale behind them, their process flow, and the analytical formulation. The modeling tools are specified and the testing methodology explained. The Section ends with an extensive results summary.

The implications of SFS on the distribution centers operations and stock targets are summarized in Section 5. A utilization analysis is developed to evaluate how the different routing models may impact the DCs day-to-day operations. The section ends by calculating a potential reduction in working capital that can be obtained as a consequence of the new ecosystem, clearly describing the approach and assumptions.

We conclude in Section 6 presenting an overview of the insights obtained throughout the paper. The section provides recommendations on specific routing models to be implemented and summarizes the conclusions and future steps considered as beneficial.

1.6 Project Objectives

The Project main goal is to provide Verizon with an innovative and accurate routing model for the DTC order fulfillment, by reviewing existing literature, conducting industry research, and testing different hypothesis. At the same time, this paper intends to provide a general solution that can be adapted to any retailer adopting an omni-channel supply chain strategy.

To consider the routing model effective, it must achieve the following objectives: decrease operational costs, reduce WC, and reduce DTC back orders, while maintaining retail stores service levels, and causing a minimal disruption to the day-to-day operations.

2 Literature and Related Research Review

Consumers are becoming increasingly demanding and expect a direct, personalized, and fast shopping experience. Modern consumers expect to be able to shop at anytime from anywhere, having full visibility to product variety, pricing information, and inventories. With a market that continuously opens to new alternatives and that transcends any kind of boundary, competition becomes fierce. To stay in business and gain customer share, businesses need to meet customer demands, so they are seeking for ways to improve customer experience without engaging in major investments in fulfillment costs and WC.

Companies' supply chains have evolved throughout the years becoming more flexible to adjust to customer preferences. This section reviews the last trends in supply chains flexibility and adaptability, citing market research, papers and examples from experienced retailers.

2.1 E-commerce

E-commerce is growing year over year with a fast pace, being one of the main drivers to supply chain evolutions. According to the National Retail Federation [4], the U.S. total retail growth for year 2017 is expected to be 3.7-4.2%, while the online retail is expected to grow by 8-12%. This means that while in-person sales keep growing, finding ways to capitalize on rising e-commerce, like through the addition of new channels or hybrid omnichannel functionality, becomes increasingly important. In addition, with the growing usage of mobile devices, online shopping is expected to keep growing in the next several years.

Moreover, a market research conducted by Deloitte explains that consumer electronics and fashion are the most important segments of growth in online retail [5]. Verizon Wireless' products fall into the consumer electronics devices category, making it essential for the company to align its strategy to satisfy this demand pattern across channels.

2.2 Omni Channel Supply Chain Strategy

Omni-channel retail (i.e. combining brick-and-mortar stores and e-tailing) is the future of e-commerce. This requires business to redesign their strategical approach when considering the demand across channels. One single channel is not enough to meet the demands of increasingly strict customers. The challenge becomes integrating among the channels to leverage the benefits of each one, without incurring massive capital investments or fulfillment expenses.

E-commerce has for many years been considered an independent channel that specific companies used as their business model by competing using less capital intensive structure. Considering the e-commerce exponential growth, traditional retailers found themselves forced to incorporate an online channel to satisfy their customers' needs. An evolved multi-channel operation required businesses to keep assets independently, including holding inventory separately and owning dedicated distribution centers to fulfill orders.

Omni-channels start as a progression of multi-channel, integrating the online and brick-and-mortar operations together to achieve several benefits and building a more efficient and responsive supply chain. Figure 2 shows this evolution of supply chains [5].

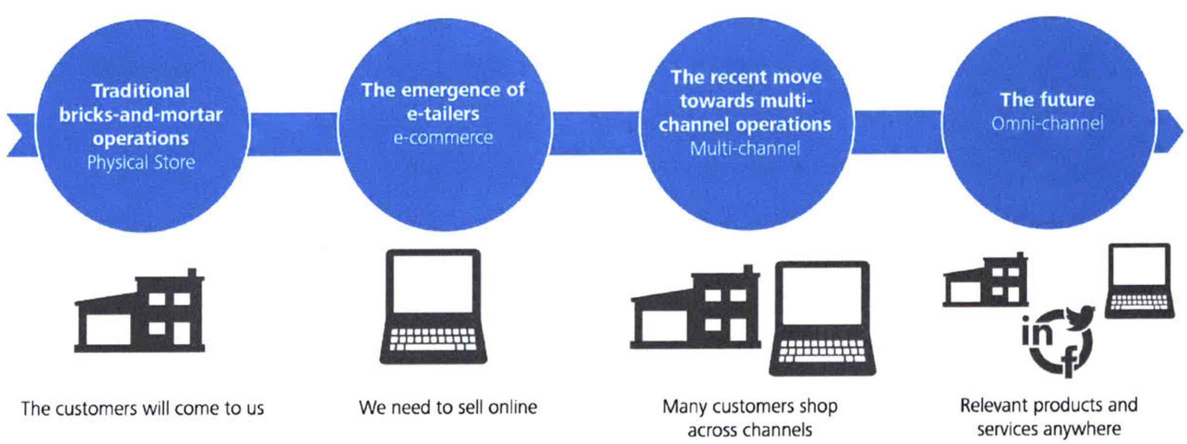

Figure 2 – Supply chains evolution

2.3 Ship from Store

Large retailers are exploring new paths to fulfill customer orders leveraging their inventory footprint. Delivering orders to a customer's home can become very costly, so using inventory already placed close to the customer can help achieve major savings.

Ship from store connects demand with inventory in a flexible way. Companies can decide between inventories at different locations to optimize fulfillment costs and inventory management. If we consider a retailer holding inventory at a DC to specifically fulfill online demand, the WC is optimized for a specific channel. Also, because DCs are sometime remotely located from the end customer, the fulfillment costs (and delivery times) are very high. This retailer, using SFS can decrease these costs by sending a product to a customer from a closer location while pooling the inventory across channels.

The outcome can be a major benefit for retailers that are already holding a massive store footprint across the country. According to a market research conducted by Fortna, several traditional retailers such as Macy's, Toys"R"Us, and Ann Taylor are implementing a ship from store strategy to fulfill direct-to-consumer demand [6]. The report cites an example regarding the apparel retailer, Ann Taylor, that "has made over 300 of their stores into distribution nodes. Their customers can order products no matter where this are. As a result, Ann Taylor was able to increase sales and gross margins by not having to mark down slow-selling items in one store that might sell at full price in another".

2.3.1 Main Benefits

The benefits that an SFS strategy entails are various. A study by the market research firm Enspire Commerce summarizes as follows [7]:

a) Competitive advantage: retailers allowing stores to ship orders are better positioned to meet customers' demand providing a more responsive shopping experience.

b) Lower costs: "Shipping from stores allow retailers to fulfill more online orders from inventory they already have, and ideally, from near the point of demand, driving down fulfillment cost".

c) Leverage aging inventory: retailers reduce in-store markdowns, and can use aging (or obsolete) inventory through a specific channel.

d) Unified inventory pools: by integrating the different fulfillment channels companies can eliminate the need to hold separate pools of inventory expected to meet the demand of each specific channel. The inventory targets can be optimized jointly.

e) Higher and faster levels of service: using store inventory to fulfill online consumers, rather than using a centralized DC, provides a quicker service, improving customer service and fostering brand loyalty.

f) IT consolidation: a single view of inventory across fulfillment channels allows to consolidate IT systems, integrating a point of sale (POS) solution to manage inventory levels.

2.4 Routing Models

Once SFS is established, a system that understands how the different inventory locations interconnect is essential to obtain the benefits summarized in Section 2.3.1. One key aspect is choosing a routing model, i.e. a model that will select which store or DC ships each customer order.

Once an order is placed, the retailer can decide how to fulfill it considering several factors such as inventory levels, outbound shipping costs, and future demand forecast across channels. To develop an efficient routing model, it is essential to consider the partial and potential visibility of demand to make better fulfillment decisions.

The first aspect of demand visibility we explore is the future implications to the inventory state, after deciding to allocate a piece of inventory from a specific location. Intuitively, retailers make fulfillment decisions by fulfilling each order the cheapest way possible given the current inventory state, without accounting for the costs implications of future fulfillment opportunities, or even losing retail sales. Incorporating these implications into the routing model can result in to major cost savings for the company [8].

Another important aspect of demand visibility to be considered is making order fulfillment decisions optimizing for a group of orders together. The approach of assigning orders to DC based on just the individual customer information is necessarily myopic because it cannot account for any subsequent customer orders [9]. We can construct a near-optimal model to assign a large set of customer orders after a period of time, allowing better fulfillment decisions based on new information.

Section 4 explains in detail the different routing models that were tested during the internship, optimizing the system and developing new practices. These practices were tested with extensive historical data from Verizon to achieve tailored results, but can be applied and generalized to other retailers with different supply chain specifications.

3 Methodology

Implementing such an extensive project as SFS is, constitutes a big disruption to the company's supply chain. Thus, Verizon has decided to implement it gradually in phases that will ensure that the transition happens smoothly across the involved teams. Each phase will have a specific objective and methodology as explained below.

This section explains the natural evolution of SFS among the phases, brings into attention future opportunities, and considers disruptions that are direct implications of SFS.

3.1 Phase I

The main objective of phase I is to engage the stores with SFS, while decreasing WC. Over the years retail stores have acquired inventory of products for which demand did not materialize, and which can be considered as obsolete or excess inventory. Phase I allows to ship from the stores only the SKUs in "excess".

Figure 3 – Phase I logics

Verizon replenishes its retail stores with devices and accessories based on expected demand. The logic behind the replenishments indicates that each store needs to have enough product to satisfy the safety stock plus the daily forecasted sales per SKU. Usually, demand materializes per the forecast, but this is not always the case. In the scenario of consistently lower demand than expected, an inventory build-up will occur for a specific store-SKU combination. In some cases, these SKUs with excess inventory are not in demand at retail stores, but some demand is seen on the DTC channel. Thus, this inventory can be used by SFS to fulfill some of the DTC orders while decreasing inventory at the stores.

For instance, suppose we define "Excess inventory" by a threshold of 10 day of supply (DOS) over the safety stock:

$n - Planning\ horizon = 10\ days\ (2\ calendar\ weeks)$

$k - minimum\ threshold\ for\ service\ level\ purposes = 1\ unit$

The units in excess inventory at each SKU-store combination is defined as the positive part of the RHS of equation 1 (i.e. if on hand is lower than the threshold, then excess inventory equals zero).

$$\boldsymbol{Excess_{St,SKU} = max\,\{OnHand_{St,SKU} - Threshold_{St,SKU}, 0]} \quad \textbf{(1)}$$

To define the threshold, we look at two specific scenarios. Equation 2 refers to the case of unforeseen demand for a specific SKU at a specific location. In this case the threshold is given by the minimum quantity ($MinQty_{St,SKU}$), which is fixed by the marketing department as a minimum requirement that stores must hold in inventory for branding purposes.

If $\sum_{t=1}^{n} Forecast_{St,SKU,t} + SafetyStock_{St,SKU} = 0$

$$Threshold_{St,SKU} = MinQty_{St,SKU} \qquad (2)$$

Equation 3 defines the threshold when demand is expected (a much more likely scenario). In this case, the threshold is defined as the larger between $MinQty_{St,\,SKU}$ and the sum of ten days of demand, plus the defined safety stock per store per SKU, and the minimum quantity for service level purposes (i.e., 'k' is fixed by the Marketing department). The level of ten days of demand is set arbitrarily to identify excess inventory that will not see demand in the immediate future, e.g., very unlikely it will be needed in the next few weeks.

If $\sum_{t=1}^{n} Forecast_{St,\,SKU,t} + SafetyStock_{St,\,SKU} > 0$

$$Threshold_{St,SKU} = Max(\, MinQty_{St,\,SKU},\, (\sum_{t=1}^{n} Forecast_{St,\,SKU,t} + SafetyStock_{St,\,SKU} + k)) \qquad (3)$$

$$SafetyStock_{St,\,SKU} = \left\lceil AdjFactor_{St,\,SKU} * \frac{\sum_{t=1}^{n} Forecast_{St,\,SKU,t}}{n} * 5 \right\rceil$$

The main economical objective of Phase I is to decrease WC. Analyzing the opportunity to use the excess inventory to fulfill DTC demand shows that in a period of 10 weeks the overall excess inventory can be reduced by more than 40%.

This result, is based on historical data over a period of five and a half months (24 weeks). The actual DTC demand was matched with the historical levels of inventory (and excess inventory) at the stores. By simulating the demand, and having calculated the excess inventory at the stores, we can replicate when excess inventory is available (using equations 1-3) and could be used to fulfill DTC demand.

Every time a DTC order triggers a store to fulfill an order from excess inventory, the overall simulated inventory level is reduced as this item will not be replenished by the DC. Given the volume of orders over a period of 24 weeks we can obtain the graph in Figure 4. As we could expect, the excess inventory decreases rapidly over a course of the first few weeks (10 weeks), and is then relatively flat thereafter.

The blue line in the graph of Figure 4 represents the excess inventory for all the SKUs that are held at the store, while the red refers only to the SKUs that are eligible through the DTC channel.

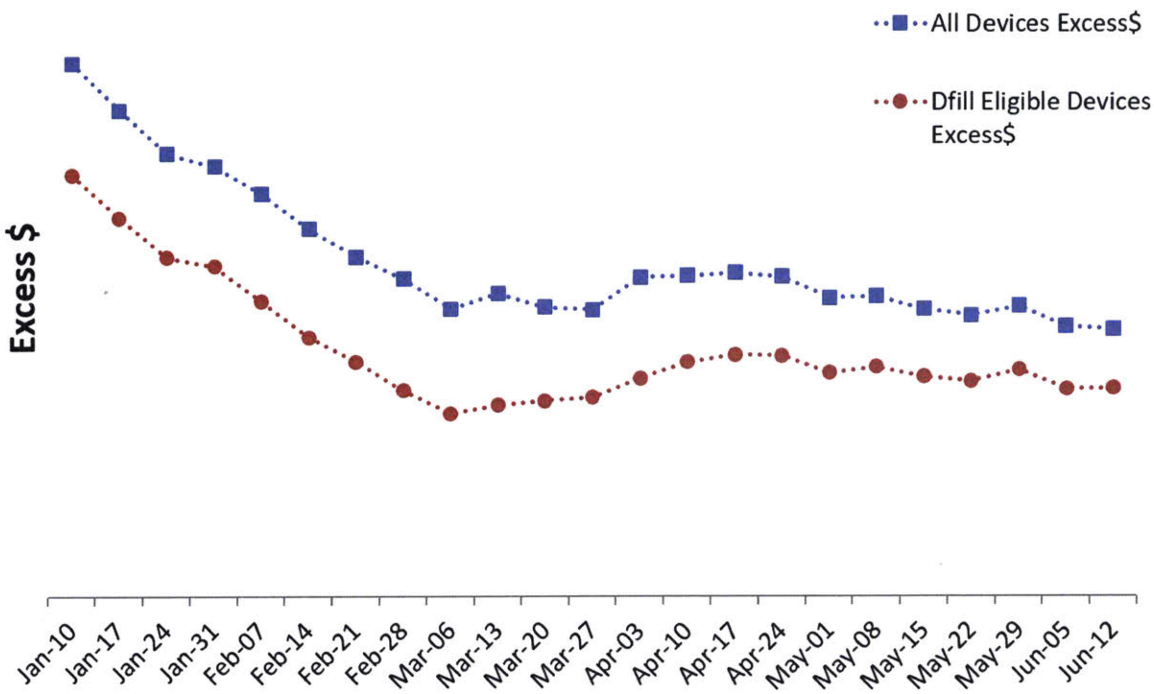

Figure 4 – Excess inventory depletion

3.2 Phase II

The main objective of Phase II is to enhance customer service levels (SL) through the DTC channel. The increase in SL is achieved by providing a backup for the DCs when these are out of inventory. Before SFS is implemented, DCs are exclusively responsible to fulfill DTC orders. If the DCs are stocked out in a specific SKU when an order is placed, the order will be automatically placed as back order. This ignores the option that a regional store may have the requested SKU in inventory and can fulfill it on time. In Phase II we will allow stores to fulfill a DTC order when there is no inventory for a requested SKU in any of the DCs. Figure 5 shows how Phase I & II work together in the same environment.

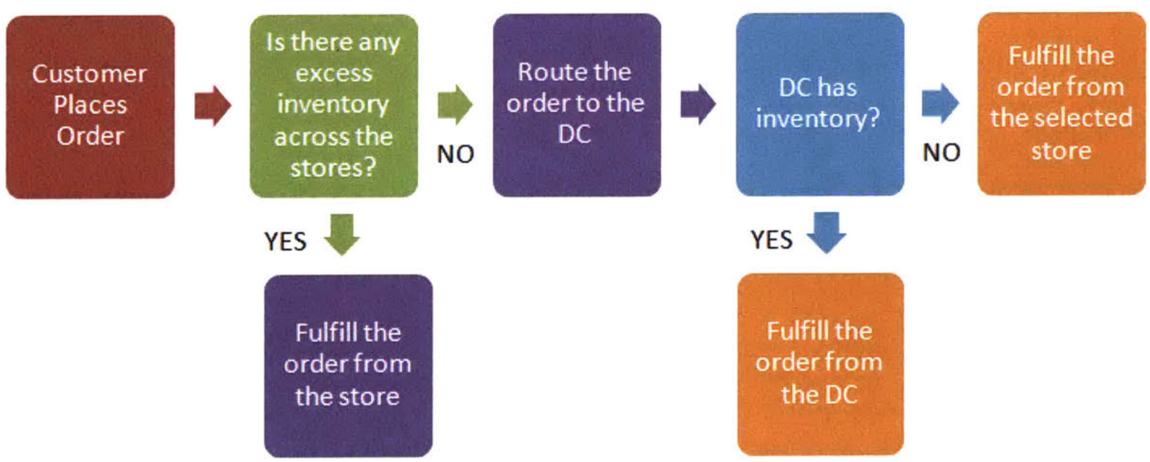

Figure 5 – Phase I and II logics combined

Based on historical data, enabling Phase II will allow Verizon to fulfill a significant number of orders yearly, which had been backordered otherwise. The benefits of being able to fulfill these orders for the company are several and include: decrease the number of new customers that may go to a competitor, reduce lost service revenues for the delayed orders, and obtain previously lost margins on accessories that could not be sold.

From the historic data for these orders, 43% are new customers. Verizon's Finance department estimates the net present value of a customer, and using this number we can assess the impact of losing one customer. Using Marketing assumptions on how many of the orders would walk away and go to a competitor, we can estimate the potential benefits of serving these customers on time, rather than putting them on backorder.

To calculate the lost margins on accessories not being purchased, we assume that customers that do not find an accessory on Verizon's inventory will go and find it with a competitor. To estimate the lost service revenue from a delayed order, we used the average number of days in delay (5 days) and the average service margin per day per customer.

Adding all these together there would be a tangible economic benefit of several millions of dollars a year in additional profits. We need to also be aware of the more intangible impact in terms of the ability to fulfill orders that the company provides its customers.

3.3 Phase III

Phase III is the ultimate omni-channel state in which Verizon will be able to fully benefit from its store footprint and inventory availability. This last stage is intended to route each single order that arrive through the DTC channel. There are no specific constraints for when stores are enabled to fulfill DTC orders in terms of DC stock outs or excess inventory, as in phases I and II. The order routing in this case attempts to find the optimal solution for given inventory states, demand forecast, fulfillment costs, and shipping constraints. Phase III will choose the best cost-effective solution while ensuring maximum customer satisfaction.

In the following sub-sections, we explain and analyze the benefits of implementing Phase III.

3.3.1 Fulfillment Costs

By implementing Phase III, fulfillment costs can be significantly decreased, using Verizon's stores footprint. Fulfillment costs are basically composed of transportation and handling:

a) Transportation costs: directly affected by using inventory that is closer to the customer. For instance, an order that had to be shipped by air from a DC in the past, can now be switched to ground by fulfilling it from a near-by store.
b) Handling costs: considers the fee per unit output that Verizon pays the third-party logistics companies that manage the DCs. Shifting a portion of the orders fulfillment to the stores decreases the overall output from the DCs, and in consequence the total handling cost declines. An important factor to consider is any associated incremental or overhead costs with the handling of DTC orders at the stores. The model assumes five orders per store per day as the maximum cap for which, no overhead cost is incurred. In other words, the operating assumption is that the existing staff attending the retail store will be able to handle up to five orders per day without compromising the day-to-

day operations and without incurring any additional operating costs. Section 4.8.2 explains and explores this parameter in assumption in detail.

3.3.2 Working Capital

Phase III will allow working capital to be decreased in two respects:

a) Avoid inventory build-up: as we have seen in Section 3.1, a large amount of excess inventory has built up over time because of SKU that became obsolete for the retail channel. Using the algorithms in Section 4 to shift DTC orders to stores when considered appropriate, will reduce this build up, catching the excess inventory at an early stage.
b) Inventory targets: before SFS the DC inventory targets per SKU were calculated based on the average throughputs seen by the DC. Using the average demand, we can calculate the safety stock and estimate the inventory target (more details in Section 5). As Section 3.3.1 explains the average throughput per SKU at the DCs will decrease and consequently the inventory targets per SKU will also be reduced.

3.3.3 Service Levels

The benefits in fulfillment costs and WC cannot be considered sufficient if SLs are compromised. Any routing model must take into consideration the impact to the customer trust on the company being able to deliver a product when needed. In order to do so, we need to consider both the retail stores and the DTC service levels.

a) Retail stores: one of the main KPI with which Verizon Wireless measures its performance is by daily tracking the service levels per SKU at the store level. For instance, stocking out at one location will negatively impact the SL of a specific SKU. For each SKU, the company defines a desired SL, usually 95% for fast-moving SKUs, and 80% for low-moving SKUs. Supply chain planning is responsible to monitor and satisfy the per-SKU service level. To prevent damaging any crucial KPI, a probabilistic SL calculation was incorporated to the routing models as Section 4.3.2 explains.

b) Direct-to-consumer: phase III maintains the benefits in Section 3.2 regarding the DTC service levels by making all the inventory available to the customers at any given moment.

3.4 DC Operations

While the benefits of implementing SFS are various and extremely appealing, we must consider other implications that may disrupt the way Verizon Wireless supply chain operates before SFS is fully implemented.

Some of the models developed will potentially change the demand flow to the DCs, by decreasing the general average demand, and by shifting when the demand arrives at the DC. Considering that DTC orders and store replenishments are handled in two independent lines at the DCs, decreasing the number of orders arriving or their arrival rate, will impact the DCs inventory targets and the utilization:

a) <u>Inventory targets</u>: the per SKU safety stock is calculated by the average demand the DCs "see" in each channel. If a portion of the DTC demand is shifted to the stores, the overall DC DTC throughput will decrease, so the calculated safety stock for the DTC channel will decrease (see Section 5.2).

b) <u>Utilization</u>: In Section 4 we look at the different routing options. One of the alternatives is to batch orders, meaning that we hold the orders for a period of time before releasing them to the DCs or stores. If we do so, the arrival rate of orders to the DCs will change and this affects its utilization (see Section 5.2).

3.5 Future Opportunities

While working on optimizing the order routing to Verizon's customers, some future opportunities came up as the natural path to explore given the new fulfillment map.

As Section 3.3.2 briefly explains, significant achievements can be obtained in terms of safety stock at the DCs considering the new demand pattern. However, a greater benefit can be found if we analyze as well the target cycle stock per SKU. At this point, average demand used for the target cycle stock will not be affected until SFS is implemented. Once implemented and consistent results achieved, in terms of what portion of demand can be shifted to the stores, cycle stock targets can be revised.

The cycle stock targets are not part of this research, but the author actively recommends to explore this in the future as a potential source of major interest for the company.

4 Phase III – Towards the Optimal Solution

This section navigates each one of the alternative routing models. It introduces the process flow and the relevant formulations for each proposed solution. At the end of the section, the final results are compared against each other. Implementation challenges, a crucial factor to be considered, are presented in Section 6 together with final recommendations.

4.1 Overview

This paper focuses on designing an optimal and effective routing model for the complex environment SFS puts in front of Verizon's supply chain. In order to develop such a routing model, several designs by the author were studied, using the literature review in Section 2 as assistance when needed.

By optimal solution we refer to the best outcome that can be achieved over different periods of time (1-hr optimization, 2-hr, daily, etc.). The best outcome is chosen when the fulfillment costs are minimized, so this would be our objective function. In order to minimize the cost, we define as decision variables the flow of product from the different locations to the customer (i.e. what store or DC fulfills each SKU). Finally, the optimal solution will be limited by several constraints such as ensuring inventory is available, maintaining service levels, fulfilling the total DTC demand per period, not exceeding the maximum number of orders per store, and non-negativity for our decision variables.

Initial analysis and test of the alternative designs allowed us to narrow the diverse models into three that are fully explained and tested in this section. Each one of these alternatives is compared to the current state and against each other, so the first thing we need to look at is how the current routing logic works and analyze its results. Thus, the following routing scenarios are presented in this Section:

a) Current state – details how orders are allocated before SFS

b) Heuristics – implement simple business rule to route orders as they arrived (FIFO)
c) Optimization Model I – order batching per a specific time frame allowing demand visibility
d) Optimization Model II – order batching on a daily/weekly basis. Not implementable, ideal solution, but used as a benchmark for comparison proposes

Figure 6 shows a simplistic diagram of the differences between the Heuristic and the Optimization models. One way to do routing is with what is known as an online algorithm – namely, you make the fulfillment or routing decision at the instant each order is placed. Another way is to make these decisions periodically for batches of orders. For instance, you might accumulate all orders that arrive over an hour, and then at the end of the hour decide how to route each of the orders.

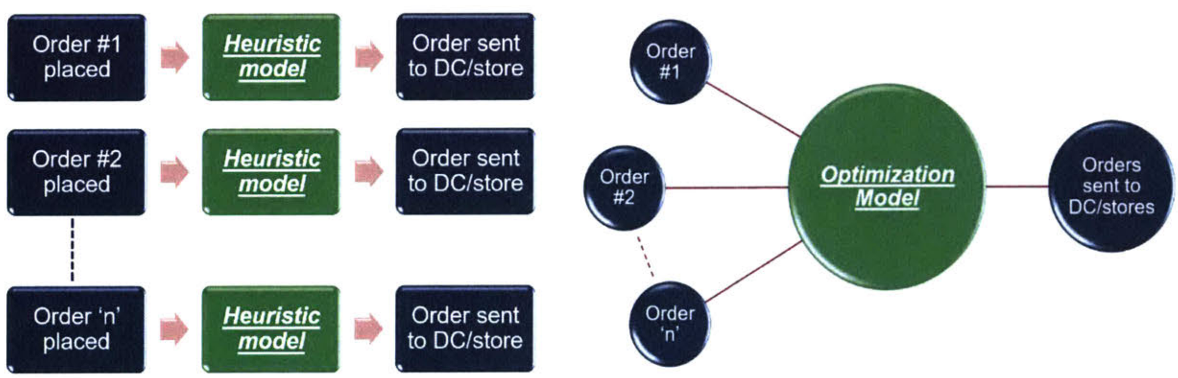

Figure 6 – Heuristic vs Optimization model

4.2 Current State

Without SFS, orders are routed based on simple business rules that dictate which DC is responsible for each order. The system works in a way that each customer zone has an assigned "Primary DC" and "Secondary DC". As the intuition indicates, once the order arrives from a specific customer zone, if the Primary DC has available inventory, the order will be directly allocated there. In the case that there is no inventory available in the Primary DC, the Secondary DC will fulfill the order if inventory is available. If none of the DCs hold stock for the ordered SKU, the order will be placed as back order.

In the case of a multi-item order, the system will work very similarly looking for all the items in the order at the Primary DC and then the Secondary DC. If there is no DC that holds all the items in the order, split orders would be considered depending on the estimated arrival time of replenishments (when next shipment should arrive to the DC), and the type of SKUs in order (devices or accessories). However, considering that the majority of DTC orders are single-item orders, and taking into account the results in Section 4.8.3 (80-85% of the benefits can be obtained by simplifying the models to single-item orders), multi-item orders are not the main focus of this paper.

Figure 7 shows a flow chart that explains the current routing process.

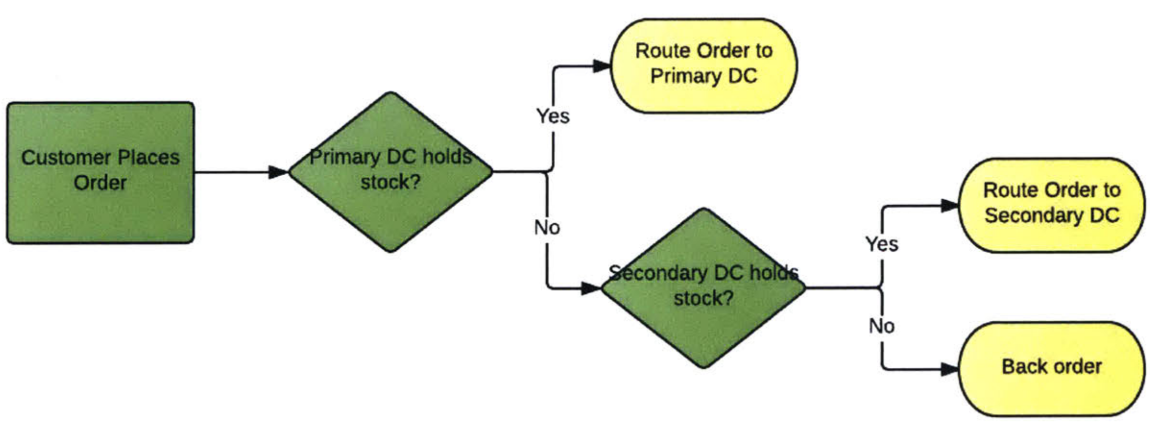

Figure 7 – Current State Process Flow

4.3 Heuristics Model

The Heuristic Model incorporates the stores as a potential outcome for the routing decision using business rules. In this model, we first filter out stores if they do not meet certain constraints such as:

a) Reached the maximum number of orders allowed per day (see Section 3.3.1)
b) Inability to meet shipping time
c) Store has all the items in the order
d) Expected retail service level drops below the minimum (to be explained in Section 4.3.2)

After these filters, the eligible stores or DCs are compared based on the total fulfillment cost, and the one with the minimal cost is selected. In case of equal fulfillment cost, the store or DC with maximum DOS per SKU is selected.

4.3.1 Process Flow

Figure 8 shows a detailed diagram of the Heuristic process flow, including potentially splitting orders for the case of multi-item orders.

With multi-item orders the model will look first into stores that have all the items ordered. If no store can meet all the constraints, splitting orders among two stores/DCs is considered finding pairs of stores/DCs holding all the items in the order. If several pairs are found, the total fulfillment costs by pair are compared to find the best routing option.

If no pair of stores was found, splitting the order is considered, weighting some of the considerations outlined in Section 4.2.

Note: in Figure 8, each time the term "store" is cited it refers to both stores and DCs as potential fulfillers of the order.

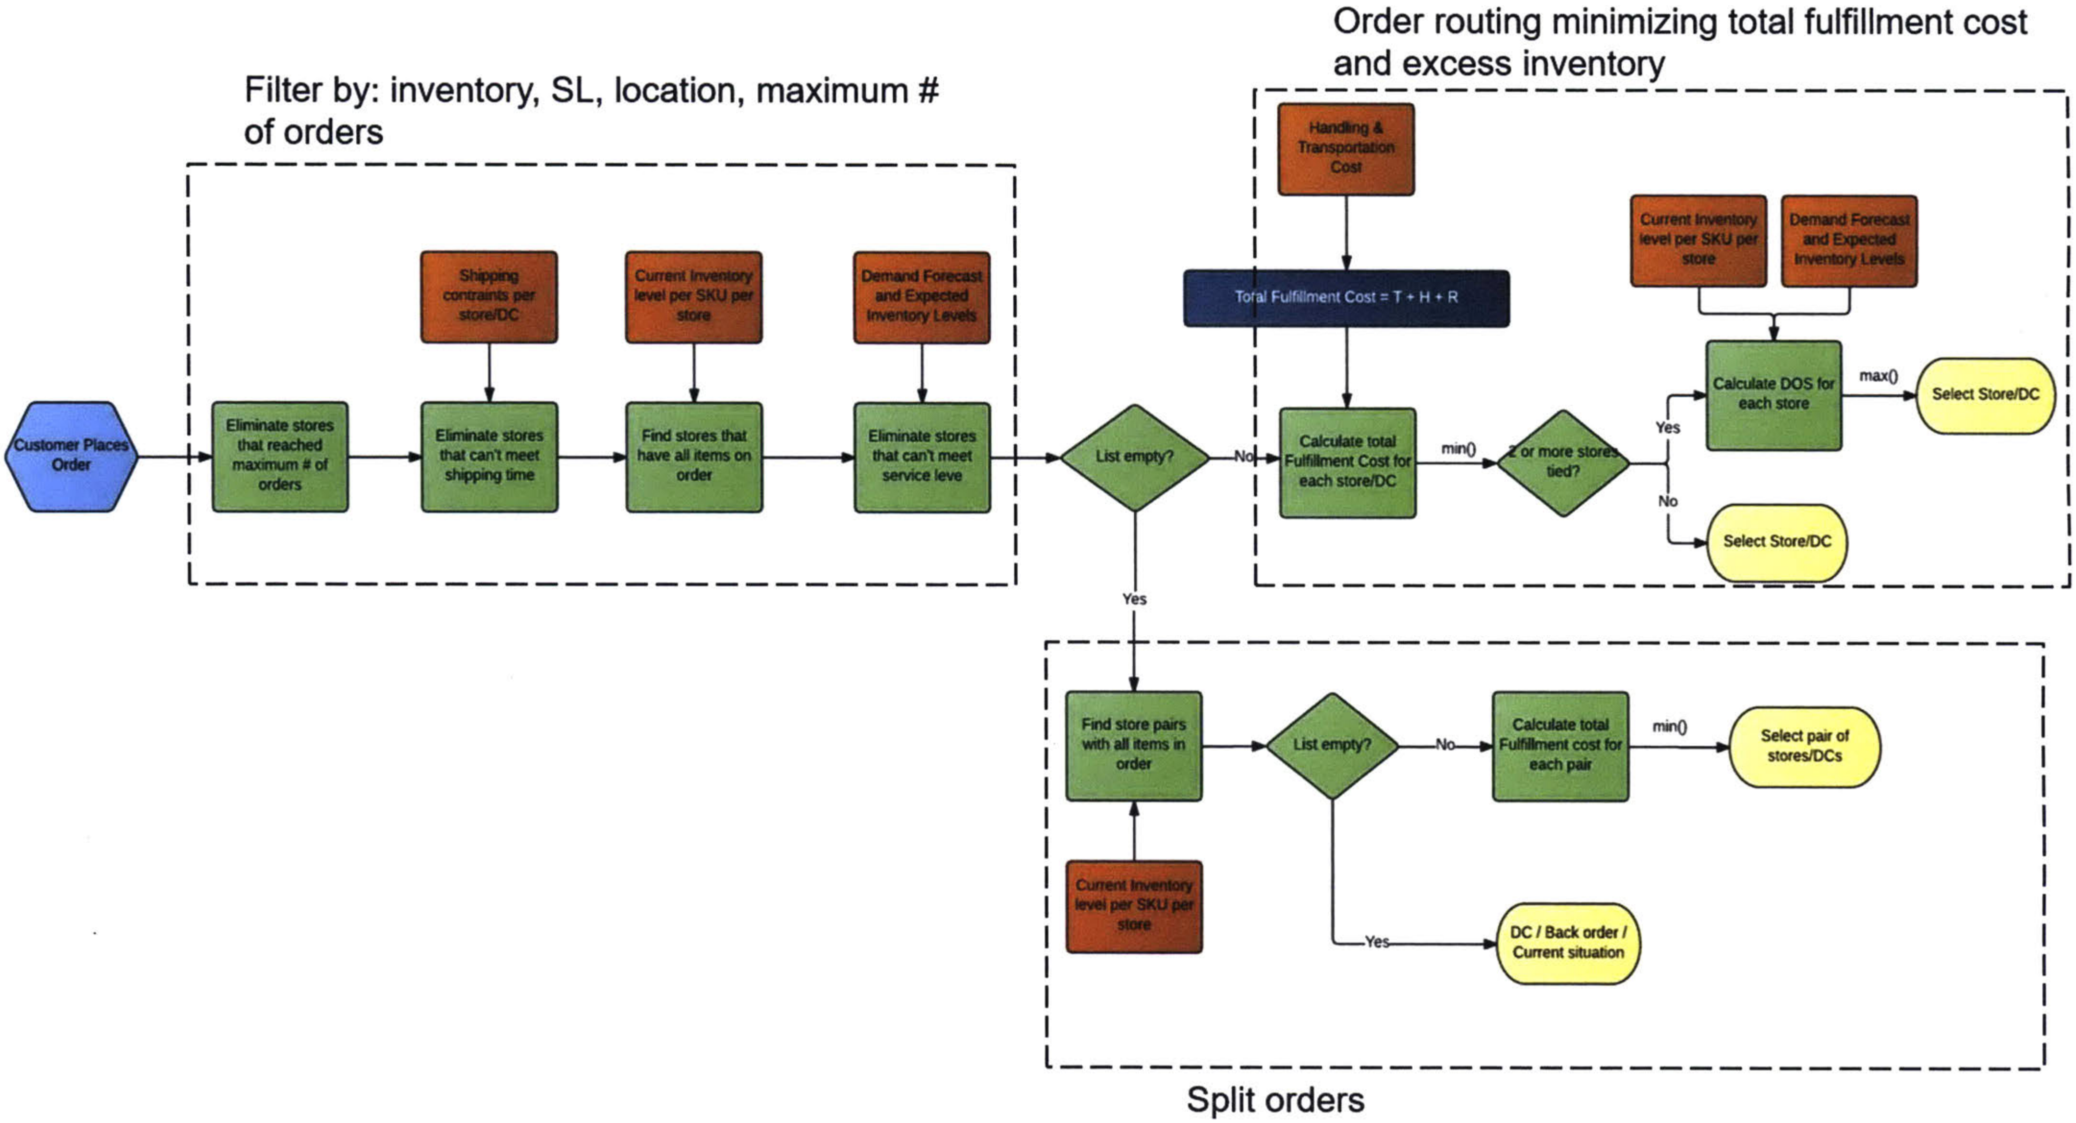

Figure 8 – Heuristic Model Process Flow

4.3.2 Service Levels

One of the main concerns when implementing SFS is the service level at the retail store. The Heuristic Model filters out stores that are or would be below the minimum defined service level if their inventory were used to serve a DTC demand. To do so, a probabilistic approach is incorporated to the model that estimates the expected level of service per SKU per store.

Considering the following definitions, we can define the service levels, on-hand inventory state, and the expected retail demand at the SKU-store level:

$i \in I - set\ of\ stores/DCs$

$k \in K - set\ of\ SKUs$

$\tau - look\ ahead\ period$

$m - day\ counter\ (0 \leq m \leq \tau)$

$\mu_{ik} -\ service\ level$

$Z_{ik}^{m} -\ expected\ inventory\ on\ hand\ of\ SKU\ k\ at\ store\ i\ at\ day\ m$

$f_{ik}^{m} - retail\ demand\ forecast\ of\ SKU\ k\ at\ store\ i\ at\ day\ m$

$b_{ik}^{m} -\ actual\ retail\ demand\ per\ SKU\ per\ location\ on\ day\ m$

Based on the definitions above the expected service level can be estimated by equation number 4:

$$\mu_{ik} = \prod_{m=0}^{\tau} \left\{ \sum_{l=0}^{Z_{ik}^{m}-2} P\left(b_{ik}^{m} = l\right) \right\} \quad (4)$$

The probabilistic approach calculates the probability of not stocking out at each specific day of the look ahead period. The look-ahead period is based on the expected lead time to replenish per store, and how quickly the supply chain reacts to replenish the "lost" item. Usually τ equals 3, as the maximum number of days that it will take to replenish the inventory at the store from a DC.

The upper limit $(Z_{ik}^{m} - 2)$ refers to how many items of a specific SKU can be sold in one day at a specific location without stocking out. The (-2) contains the item potentially fulfilled by SFS, and the minimum quantity to avoid a stock out.

The Probability within the round brackets, assumes a Poisson distribution with the demand forecast as average. Summing up the probabilities from not selling any item to the upper limit $(Z_{ik}^{m} - 2)$, constitutes the probability of not stocking out on a specific 'm' day.

In order to calculate the probability of not stocking out at any given day throughout the look ahead period, each day is considered as independent, which is an approximation. Thus, they are multiplied by each other by the product symbol.

The result for μ_{ik} is the approximate probability of not stocking out in the defined look ahead period, and it represents a conservative method to simulate the service level of a SKU at a specific location.

As an example, let's consider a specific store with the following levels of expected inventory for a specific SKU1:

	Day m=0	Day m=1	Day m=2	Day m=3
Inventory (Z)	8	7	7	8
Retail demand forecast	1	0	1	1
Replenishment expected	0	0	2	0

Given the levels of inventory and the expected demand (the average of the Poisson distribution), we can calculate the probabilities of selling either 0, 1, 2, etc. items of SKU1. For instance, the probability of selling 0 of SKU1 on day m=0 is P(b=0) with average demand of 1 (always assuming Poisson), and the result would be 0.368. If we calculate this probability for b=1,2...,6 and the sum all the probabilities, we obtain the daily service level. In this case, this is the sum of probabilities up to 6, as it represents the parameter $(Z_{ik}^{m} - 2)$, being $Z_{ik}^{m} = 8$ for day m=0. Adding all the

Poisson probabilities, we obtain that the likelihood of not stocking out at this location on day m=0, if one item of SKU1 is used to fulfill an SFS order, is 0.9994.

Repeating this calculation for each day, we obtain the daily probability of not stocking out (or the daily service level). Multiplying the daily service levels by each other, assuming independent event (conservative assumption), obtains the expected probability of not stocking out across the 3-day period. In this case we get 0.995 or service level of 99.5%.

4.4 Optimization Model I

Delaying the allocation decision for orders for a short period in order to increase demand visibility is considered in this section. By not routing the orders instantly as they arrive, we allow the opportunity to optimize among several orders together. Optimization Model I batches orders among periods such as one hour of orders, two hours, and three hours. As the batching period increases, more orders are optimized together, and theoretically a greater benefit can be obtained. However, a larger batching period corresponds to a longer delay between order arrival and when it is allocated to a store or DC, and this can become unfeasible as explained in Section 4.5.

4.4.1 Linear Program Formulation

A linear program (LP) was used to formulate the Optimization Models. This section explains in detail the formulation, including the parameters, decision variables, inputs, objective function, and constraints.

<u>Parameters</u>

$i \in I - set\ of\ stores/DCs$

$k \in K - set\ of\ SKUs$

$j \in J - set\ of\ customer\ zone$

$t - 0\ or\ T\ (before\ and\ after\ running\ the\ algorithm)$

$M_i - counter\ to\ keep\ track\ of\ the\ number\ of\ orders\ shipped\ per\ day\ per\ store$

Customer zone 'j' is designated by a geographical location (zip code), and by a service class (same day, next day, two day). So, customer 'j' might be Basking Ridge, with 2-day shipping.

When the parameter 't' is zero it means that no SFS orders were allocated yet, and when 't' equals T, SFS orders were allocated. This parameter will help us with inventory balances.

Decision Variables

$X_{ijk} -$ $single - item\ flow\ of\ SKU\ k\ from\ store\ or\ DC\ i\ to\ customer\ zone\ j$

$Y_{ijk} - multiple - item\ unsplit\ flow\ of\ SKU\ k\ from\ i\ to\ customer\ zone\ j$

$Z_{ik}^{T} - end\ inventory\ available\ for\ SFS\ of\ SKU\ k\ at\ location\ i\ at\ time\ T$

The Decision Variable "Y" applies for multi-item orders to be shipped from a specific location. This means the flow of a specific SKU 'k' from store 'i' to customer zone 'j', but only when shipped together with other items. Multi-item orders that need to be split are considered with the parameter v_k as the penalty for shipping multi-item orders in more than one package. While the parameter v_k constitutes an approximation, it accounts for split orders even when these are a very small portion of the orders. It is important to notice that for the mathematical analysis the split orders have very little to zero impact.

Inputs

$Z_{ik}^{0} -$ $initial\ SFS\ available\ inventory\ per\ SKU\ per\ location\ at\ time$ 0

$d_{jk} -$ $DTC\ demand\ for\ SKU\ k\ from\ customer\ zone\ j\ during\ time\ T$

$b_{ik}^{m} -$ $actual\ retail\ demand\ per\ SKU\ per\ location\ on\ on\ day\ m$

$f_{ik}^{m} -$ $retail\ demand\ forecast\ per\ SKU\ per\ location\ on\ day\ m$

$C_{ij} -$ $fulfillment\ cost\ from\ store\ or\ DC\ i\ to\ customer\ zone\ j$

$C_{ij} = T_{ij} + H_i + R_i$

$T_{ij} - Transportation\ cost\ from\ store\ i\ to\ customer\ zone\ j$

$H_i - Handling\ cost\ at\ store\ or\ DC\ i$

$R_i = H_{i@DC_b} + 0.05 * T_{DC\ to\ i}$

$H_{i@DC_b} -$ $bulk\ handling\ at\ assigned\ DC$

$T_{DC\ to\ i} -$ $Transportation\ cost\ for\ large\ bulk\ boxes$

$\rho_k - proportion\ of\ customers\ ordering\ SKU\ k\ with\ other\ items$

$w_k - expected\ discount\ per\ unit\ for\ sending\ \ a\ multi - item\ order\ in\ a\ single\ package$

$v_k - penalty\ per\ unit\ for\ split\ orders$

The input data is obtained from Verizon's database. Inputs such as inventory, demand forecast, sales, and handling and transportation costs, are available to pull. However, there are some input parameters that needed to be calculated:

a) Z_{ik}^0 – SFS available inventory refers to a calculated number before each run that dictates how many items can be shipped out of a SKU at a specific location without impacting service levels. Z_{ik}^0 is calculated again after each run, based on the service level equation in Section 4.3.2.

b) R_i – defined as the cost of replenishing an item at a store 'I', is calculated adding the costs of handling the specific SKU at the DC plus the cost of transporting it to the store. The parameter $H_{i@DC_b}$ means the cost of handling at the DC when the SKU is picked for bulk (retail) replenishment. In this case the handling cost per unit is significantly lower that the DTC handling cost. The parameter $T_{DC\ to\ i}$ is the transportation costs of a big replenishment box from a DC to a store. Usually this boxes contain in average 20 SKUs, so the coefficient (0.05) that multiplies accounts for the contribution of a single SKU to the transportation cost.

c) ρ_k – the portion of customers ordering multi-item orders is estimated using historical data for 3 months of demand. For each SKU ρ is the percent of its demand that occurs as part of a multi-item order. For instance, if SKU1 is ordered 10 times in a period of time, and 3 of them are as part of a multi-item order, then ρ=0.3.

d) w_k –the allocation of shipping cost to each SKU unit, when the SKU is part of a multi-item order shipped in single box. This parameter is based on shipping prices data, and weighted to obtain an estimate.

e) v_k – this parameter accounts as a penalty for shipping orders in multiple packages and from different locations. It includes weighted in the portion of orders that actually have to be split based on historical data.

Objective Function

The objective of the LP is to minimize the aggregate total fulfillment cost for all orders. In the objective function, we incorporate the discount of sending multiple-items in one single shipment. To simplify the model, splitting orders is limited to a maximum of two separated shipments.

$$\boldsymbol{min}\{\sum_{i,j,k} \boldsymbol{C_{ij} * X_{ijk}} + \sum_{i,j,k} \boldsymbol{v_k * w_k * C_{ij} * Y_{ijk}}\}$$

Constraints

$$Z_{ik}^{T} = Z_{ik}^{0} - \sum_{j} X_{ijk} - \sum_{j} Y_{ijk} \qquad \forall\, i, k \;\; (inventory\ balance)$$

$$\textstyle\sum_i X_{ijk} + \sum_i Y_{ijk} = d_{jk} \quad \forall\, j,\ k \;\; (DTC\ demand\ fulfillment)$$

$$\sum_i X_{ijk} = (1 - \rho_k) d_{jk} \qquad \forall j, k \qquad (single - item\ orders)$$

$$\sum_i Y_{ijk} = \rho_k d_{jk} \;\; (muktiple - item\ orders)$$

$$M_i + \sum_{j,k} X_{ijk} + \sum_{j,k} w_k * Y_{ijk} \le 5 \qquad \forall\, i, m \quad (maximum\ 5\ orders\ per\ store\ per\ day)$$

$$\left.\begin{array}{ll} Z_{ik}^{T} \ge 0 & \forall\, i, m, k \\ X_{ijk}, Y_{ijk} \ge 0 & \forall\, i, j, m, k \end{array}\right\} \textit{Non - negativity}$$

After each run

a) $M_i {+}{=} \sum_{j,k} X_{ijk} + \sum_{j,k} Y_{ijk} + \sum_{j,k} W_{ijk}$ - increase M_i by the total numbers of orders shipped from store I during the last period.
b) $X_{ijk} = 0; Y_{ijk} = 0; W_{ijk} = 0$ – set the decision variables to zero
c) Calculate Z_{ik}^{0} – available SFS inventory based on equation (1).

4.5 Optimization Model II

This model is an expansion of Model I since it allows to batch orders in a longer period as a few days or a week. Clearly, orders cannot be held for such long periods as a day or a week, before being routed to a store or a DC, if the company wants to stand by the shipping commitments to its customers. Thus, this model was developed with the purpose of comparing the previous models to an "ideal" solution, but not with the intention of implementing it on Verizon's systems.

4.5.1 Linear Program Formulation

The formulation in this case is very similar to the one presented in Section 4.4.1 The main difference is the way we treat the inventory levels. For instance, if we optimize for an entire week, we obtain full demand visibility and there is no need to calculate the available SFS inventory as in Section 4.4. We can allocate orders in a way that all the retail demand is already allocated, so the stock outs become less relevant.

When considering multiple days into the model, we need to introduce the expected and actual replenishment, as it has a significant impact on the inventory levels. The replenishment constraints are modeled per Verizon's replenishment policy. These are targeted to achieve a beginning-of-day stock level of the sum of the safety stock and forecasted daily sales per SKU per location.

<u>Parameters</u>

$i \in I - set\ of\ stores/DCs$

$k \in K - set\ of\ SKUs$

$j \in J - set\ of\ customer\ zone$

$\tau - number\ of\ days\ in\ period$

$m - day\ counter\ (0 \leq m \leq \tau)$

Decision Variables

$\boldsymbol{X_{ijmk} - single - item\ flow\ of\ SKU\ k\ from\ \ i\ to\ customer\ zone\ j\ on\ day\ m}$

$\boldsymbol{Y_{ijmk} -\ \ multiple - item\ unsplit\ flow\ of\ SKU\ k\ from\ i\ to\ \ j\ on\ day\ m}$

$\boldsymbol{Z_{ik}^{m} - inventory\ of\ SKU\ k\ at\ location\ i\ at\ end\ of\ day\ m}$

$\boldsymbol{R_{ik}^{m} -\ \ planned\ inventory\ replenishment\ of\ SKU\ k\ at\ location\ i\ at\ start\ of\ day\ m}$

Inputs

$Z_{ik}^{m} -\ \ inventory\ levels\ per\ SKU\ per\ location\ 3\ (m-1, m-2, m-3)\ days\ back$

$Z_{ik}^{0} -\ \ initial\ inventory\ per\ SKU\ per\ location$

$L_i -\ \ lead\ time\ (in\ days) to\ replenish\ store\ i\ from\ its\ primary\ DC$

$R_{ik}^{0} -\ \ replenishment\ on\ day\ 0\ (and\ day\ 1\ if\ L_i = 2)\ per\ SKU\ per\ location$

$d_{jk}^{m} -\ \ DFill\ demand\ forecast\ for\ SKU\ k\ from\ customer\ zone\ j\ on\ day\ m$

$b_{ik}^{m} -\ \ actual\ retail\ demand\ per\ SKU\ per\ location\ on\ day\ m$

$f_{ik}^{m} -\ \ retail\ demand\ forecast\ per\ SKU\ per\ location\ on\ day\ m$

$C_{ij} -\ \ fulfillment\ cost\ from\ store\ or\ DC\ i\ to\ customer\ zone\ j$

$SS_{ik} -\ \ Safety\ stock\ (or\ minimum\ qty)\ per\ SKU\ at\ store\ i$

$C_{ij} = T_{ij} +\ H_i +\ R_i$

$T_{ij} - Transportation\ cost\ from\ store\ i\ to\ customer\ zone\ j$

$H_i - Handling\ cost\ at\ store\ or\ DC\ i$

$R_i = H_{i@DC_b} + 0.05 * T_{DC\ to\ i}$

$H_{i@DC_b} -\ \ bulk\ handling\ at\ assigned\ DC;$

$T_{DC\ to\ i} -\ \ Transportation\ cost\ for\ large\ bulk\ boxes$

$\rho_k - proportion\ of\ customers\ ordering\ SKU\ k\ with\ other\ items$

$w_k - expected\ discount\ for\ sending\ \ a\ multi - item\ order\ in\ one\ single\ package$

$v_k - penalty\ for\ split\ orders$

Like in Section 4.4.1 input data is obtained from Verizon's database. Inputs such as inventory, demand forecast, sales, and handling and transportation costs, are available to pull directly from the databases. The calculated inputs are equivalent to Section 4.4.1.

Objective Function

The objective of the LP is to minimize the aggregate total fulfillment cost for all orders. In the objective function, we apportion to each SKU the cost of sending multiple-items in one single shipment, based on the parameter w_k. To simplify the model, splitting orders is limited to a maximum of two separated shipments.

$$\boldsymbol{min\{\sum_{i,j,m,k} C_{ij} * X_{ijmk} + \sum_{i,j,m,k} v_k * w_k * c_{ij} * Y_{ijmk}\}}$$

Constraints

$$\left.\begin{array}{l} if\ L_i = 1 \rightarrow R_{ik}^m = f_{ik}^m + f_{ik}^{m+1} + SS_{ik} - (Z_{ik}^{m-2} + R_{ik}^{m-2} + R_{ik}^{m-1} - f_{ik}^{m-1}) \\ if\ L_i = 2 \rightarrow R_{ik}^m = f_{ik}^{m-1} + f_{ik}^m + f_{ik}^{m+1} + SS_{ik} - (Z_{ik}^{m-3} + R_{ik}^{m-3} + R_{ik}^{m-2} + R_{ik}^{m-1} - f_{ik}^{m-2}) \end{array}\right\} \textit{Replenishments}$$

$$Z_{ik}^m = Z_{ik}^{m-1} + R_{ik}^m - b_{ik}^m - \sum_j X_{ijmk} - \sum_j Y_{ijmk} \quad \forall\, i,k,m \ (inventory\ balance)$$

$$\sum_i X_{ijmk} + \sum_i Y_{ijmk} = d_{jk}^m \quad \forall\, j,m,k \ (DTC\ demand\ fulfillment)$$

$$\sum_{j,k} X_{ijmk} + \sum_{j,k} w_k * Y_{ijmk} \le 5 \ (maximum\ 5\ orders\ per\ store\ per\ day)$$

$$\sum_i X_{ijmk} = (1-\rho_k) d_{jk}^m \quad \forall j,k \quad (single - item\ orders)$$

$$\sum_i Y_{ijmk} = \rho_k\, d_{jk}^m \ (muktiple - item\ oreders)$$

$$\left.\begin{array}{l} Z_{ik}^m \ge 0 \quad \forall\, i,m,k \\ X_{ijmk}, Y_{ijmk} \ge 0 \quad \forall\, i,j,m,k \\ R_{ik}^m \ge 0 \quad \forall\, i,m,k \end{array}\right\} \textit{Non - negativity}$$

4.6 Modeling Tools

Visual Basic was used to simulate the behavior of the supply chain using the Heuristic Model as the DTC routing model. This was also the tool used to simulate the current state of the routing process (before SFS).

CPLEX was the preferred software to model and run the Optimization Models. Its capabilities allowed us to effectively obtain results for a vast number of variables and parameters, as required by these complex models.

Other software such as R, Python, and Excel were used to formatting data, and analyzing results.

4.7 Testing Methodology

The first step to start performing the simulations was to acquire relevant, significant, relatively extensive, and non-biased data. In order to achieve this, historical data was used to simulate demand and the order patterns. To make sure the dataset is not biased and extensive enough, several months of data were analyzed, choosing months where no significant disruptions or promotional events were taking place (i.e. a new phone launch).

While we want the models to work in unstable situations, there are some special considerations. For example, during a new phone launch the company might prioritize retail stores to have inventory, not allowing them to ship the new device from the stores. Moreover, during these launches, the inventory for this SKU and their service levels are extremely compromised, so the models will not consider SFS as an appropriate alternative.

The data used corresponds to a two-month period. The dataset included real orders that can choose among +1,000 different SKU and can be shipped from multiple hundreds of stores and several DCs. On average, thousands of DTC orders are placed every day, and the majority of them are single-item orders. As well, most of the orders are either single-item or two-item orders.

Once the data were obtained, and its consistency analyzed, the models were run for different intervals and combinations. The final results on Section 4.8 were extrapolated from the dataset to an entire year. The scaling was done assuming the monthly benefits are even across the year. To stay in the conservative side, and to include one month of disruptions (launches, promotions, etc.), we considered the total yearly benefits on an 11-month basis.

4.8 Results Summary

This Section summarizes the results obtained through the simulations and compares the alternative routing models. The results compare the current state before SFS implementation, to the Heuristics and the Optimization Models based mainly on the total fulfillment cost to supply the DTC demand across one year. A sensitivity analysis is also performed, increasing the maximum number of orders allowed per day per store. As well, other considerations are explained in terms of retail operations disruption and the benefits of simplifying the model to single-item orders.

4.8.1 Fulfillment Costs and Potential Savings

The main metric by which the quality of each routing model is measured is the total fulfillment costs and the potential savings. These include the cost of handling, shipping and replenishing items according to the demand patterns and forecast. Any small relative saving in fulfillment costs is very valuable given the fact that currently multi-million dollars are spent in fulfilling the DTC demand each year. Figure 9 shows the yearly spending in fulfillment costs by scenario, comparing the current state to the Heuristic Model, and different batching periods of the Optimization Models. The total spending decreases as we add complexity and visibility to the models. Given the results in Figure 9, if weekly gives the best possible scenario, then the Heuristic Model achieves 60% of the maximum theoretical benefit. The short-term LP models achieve on average 70% of the weekly batching alternative.

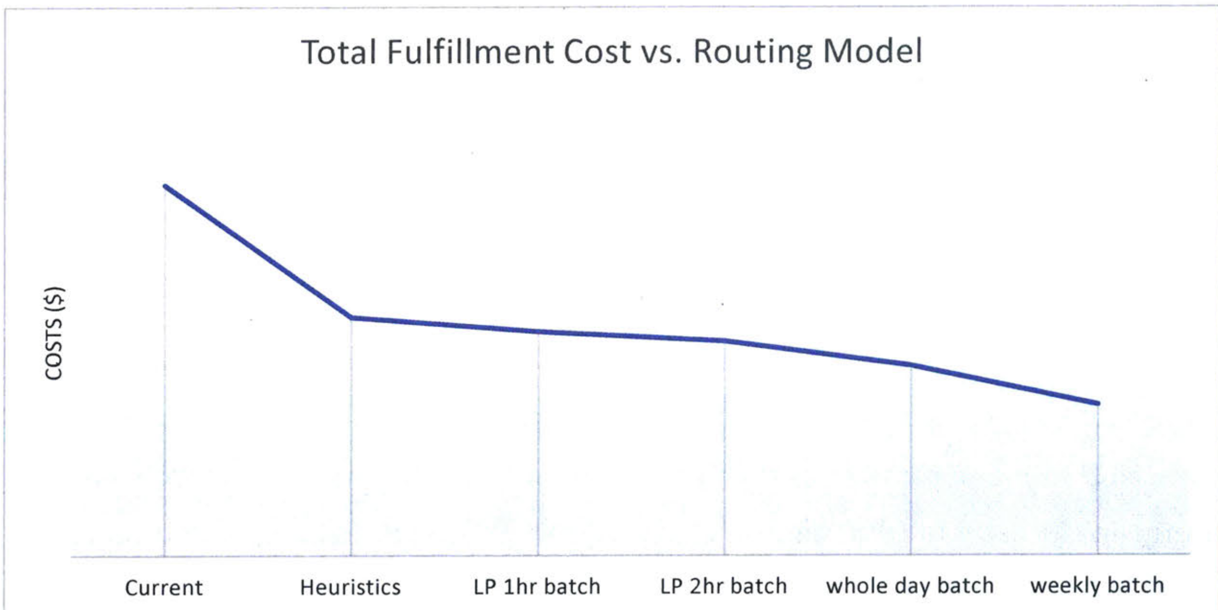

Figure 9 – Total Fulfillment costs

Note: the graph shown in Figure 9 is masked for non-disclosure purposes and it is not fully scaled. It merely represents an illustration of the cost savings, while the relative savings are explained below.

Considering the total costs, the potential savings from the current state per alternative can be calculated. Figure 10 introduces the total savings as a function of the routing alternative. Implementing the Heuristic Model yields an expected yearly saving of 3.6% in total fulfillment costs. Moreover, the feasible Optimization Model I obtains an incremental savings of 0.4-0.7%, depending on the batching period. For the "ideal" solution of long-term batching (one-week) the Optimization Model II results in a 6.2% savings in yearly fulfillment costs, but this is not practically achievable, as is explained in Section 4.5.

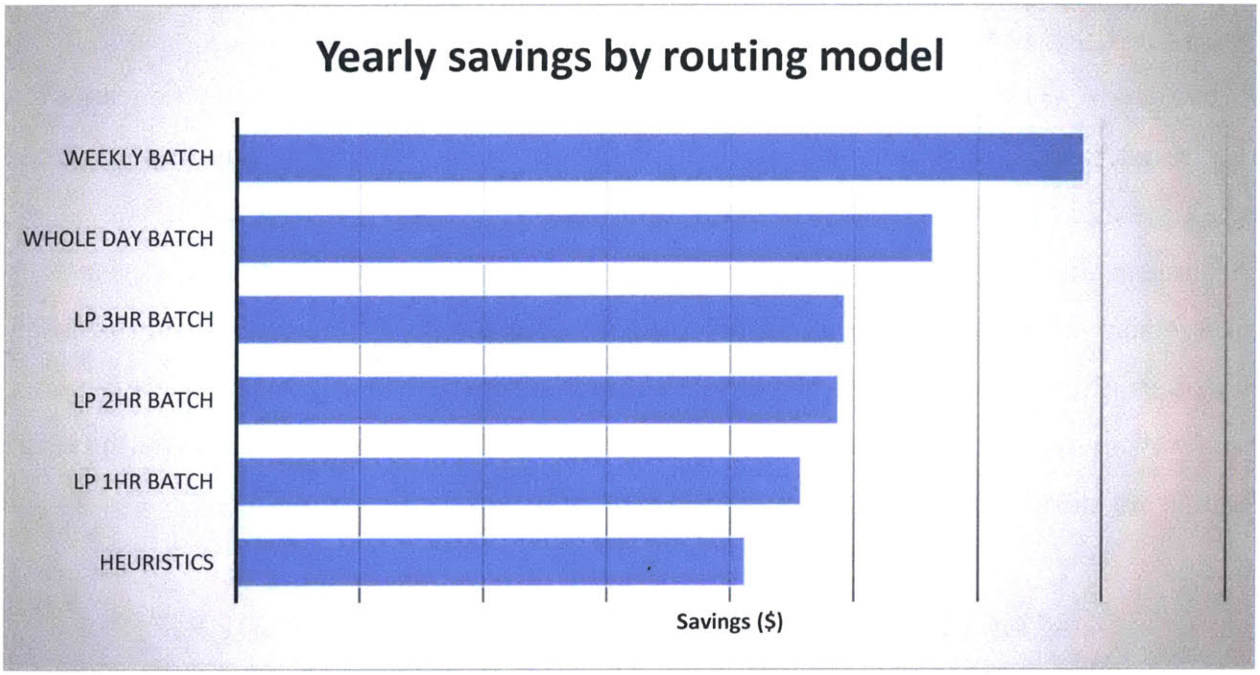

Figure 10 – Total savings by routing model

Figure 10 also proves the capabilities of the Heuristic Model if we compare it to the feasible Optimization Models (short-term 1-hr, 2-hr, etc.). By implementing a simple logic based on well-defined business rules, 85-90% of the benefit can be obtained. This is a very good approximation of a scenario in which the demand visibility is greater and the routing optimized accordingly.

4.8.2 Orders per Store per Day

The results that are presented above are constrained by allowing a maximum of five orders fulfilled per store per day. In order to understand the potential existing flexibility in Verizon's supply chain, the models were tested with some variations to this constraint. This sensitivity analysis is presented in Figure 11. The impacts of relaxing this tight constraint are very significant. For instance, increasing the maximum number of orders allowed per day to ten allows the benefits to increase, on average, by 50%, depending on the routing alternative. As the maximum number of orders allowed gets higher, the benefits grow almost linearly until 20 orders per day, with a decreasing growth rate between 20 and 40, and finally converge, having almost no significant marginal impact at the subsequent levels.

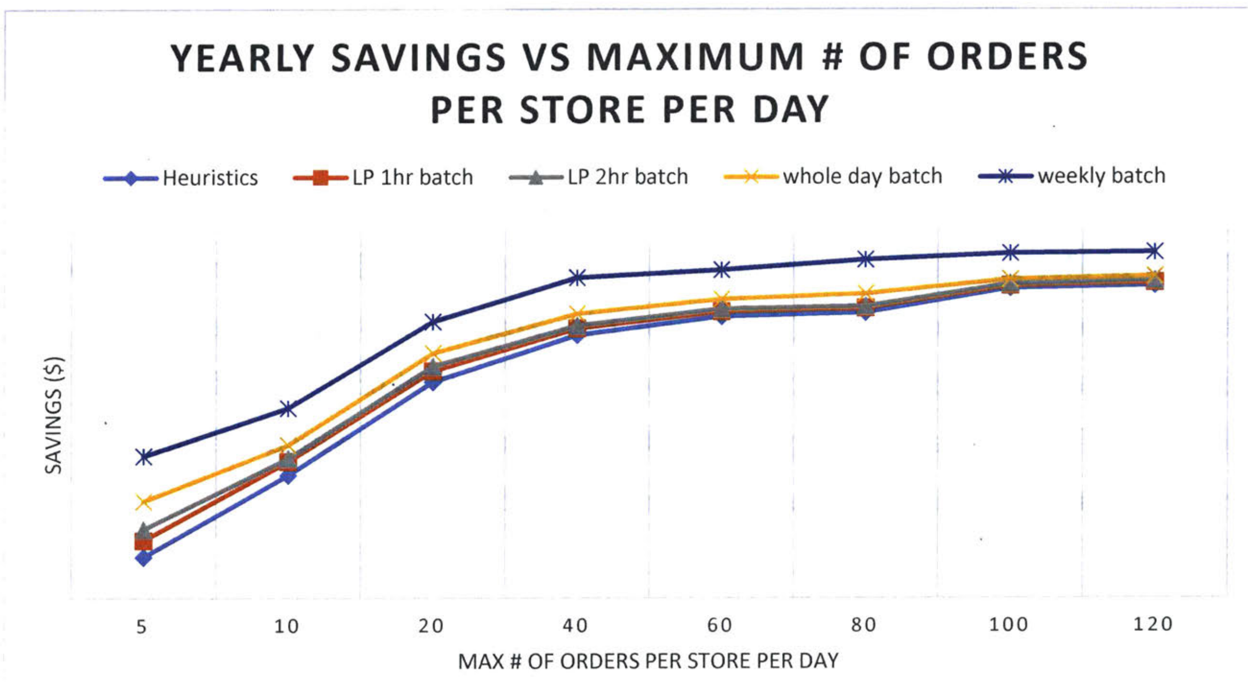

Figure 11 – Sensitivity Analysis depending on # of orders per store per day

Increasing the maximum number of orders allowed per store per day allows the routing models to obtain greater benefits. As an example, if we allow the stores to fulfill 20 orders per day, and use the Heuristics as the routing model, the yearly savings increase to an 8.0%, more than double the result obtained when only five orders were allowed. Having this significant benefit available, it is important to grasp the implications of implementing such a change.

Section 3.3.1, Part b. explains the logic behind the maximum number of orders per store per day. It is crucial to recognize that incurring an additional workload at the retail stores has a negative impact both in the financials and the day-to-day operations, which can easily exceed the benefits in fulfillment costs. To assess the operational impact of allowing the stores to ship more orders, we can look at how many orders were in fact shipped out per store during the simulation. In other words, although the stores are allowed to ship, for example, 20 orders per day, in fact they are only shipping a fraction of this number. Figure 12 presents the results of this analysis, including the standard deviation of the number of orders shipped across the stores.

Figure 12 – Actual # of orders shipped per day

Figure 12 shows that extending the maximum number of orders per store does not have a large impact on overhead costs on average, as the average number of orders per store remains relatively low. To illustrate, when the stores are allowed to ship out up to 40 orders per day, in fact they ship on average only eight orders.

If we look at the variation of the number of orders shipped across the stores, we can see that the standard deviation is between 30-50% of the average. Analyzing the patterns among the stores, we can identify that the largest stores (usually destination stores in big cities) are commonly closer to the upper limit. This result does not contradict the takeaways from this analysis, as this stores have more staff and probably could handle the order load. In addition, in the case that in a specific day a store cannot handle the number of orders, these can always be redirected to the DCs for fulfillment. Receiving an additional small number of orders will not be an issue for the DC to handle, and the economic benefits will not be significantly impacted.

4.8.3 Single-item vs Multiple-item orders

The impact of limiting the models to only single-item orders was also tested during the analysis. Per request of management, and as a potential solution to simplify implementation, in this section we measure the portion of the benefits that are obtained if stores are only allowed to ship single-item orders. In this case, multiple-item orders are shipped by the current routing model from Section 4.2.

Figure 13 compares the outcome in two specific and very relevant scenarios of a maximum of 5 or 20 orders allowed per day per store, also with a comparison between the Heuristic Model and the Optimization Model I.

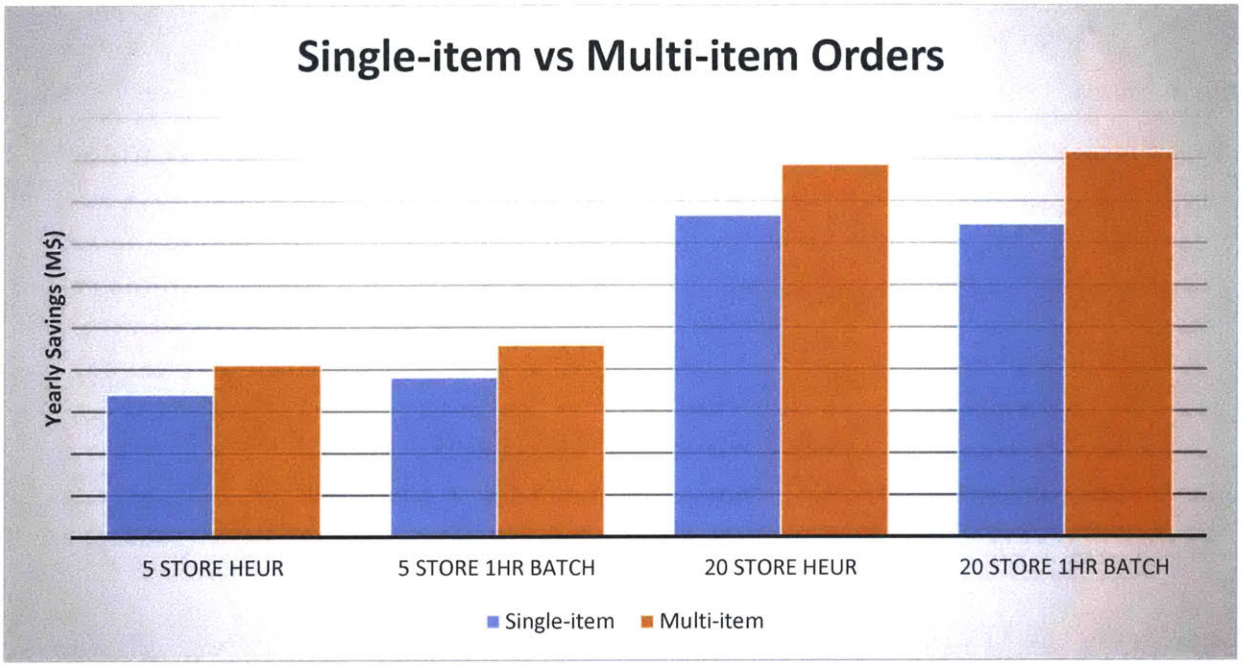

Figure 13 – single-item orders only

Enabling stores for only single-item orders yields 80-85% of the benefits of the full-scale model in which also multiple-item orders are allowed. This is a promising result that argues in favor of implementing a simplified model and can be considered as an effective alternative. The disruption for IT organization and the investment regenerating the automated systems will significantly decrease, while most of the benefit is still achievable.

5 DC Operations

The economic benefits by themselves are not sufficient to make the business case for implementation. We need to consider the challenges in implementation as well as the investments required by the company to obtain those benefits.

Once a routing model is approved to be implemented several changes will need to occur on the overall supply chain. The main stakeholders affected will be:

- Information Technology – implements the new model in the system
- Distribution Centers – disruption of the order arrival pattern
- Stores – adopting the new responsibility of shipping orders

Considering the IT organization oversees incorporating the routing model to the automated systems, one area that is crucial to evaluate when deciding between models is the "available to promise" (ATP) check. The ATP check confirms that a piece of inventory is certainly available to purchase and it is performed before the customer checks out. This test can take up to several seconds to perform for a specific SKU-store combination. As Section 6.2 explains, the ATP check is a tactical challenge to be solved especially for the Heuristic Model, where the order is routed simultaneously with the purchase (in contrast to the Optimization Models where the orders are routed behind the scenes).

While the changes to the stores are not trivial, these are not considered in this paper, as it constitutes a challenge of the entire SFS project and not specific to the routing model. Some considerations are mentioned on Section 4.8.2. regarding the potential overhead costs as shipping becomes a responsibility of the stores' workforce.

This section focuses on analyzing and explaining benefits and implications of SFS for the Distribution Centers mainly in two aspects: utilization and inventory targets.

5.1 DC Utilization Analysis

In order to decide among the different routing models, it is crucial to understand the implications these may have to the warehouse day-to-day operations and how they may disrupt (or not) it. As the warehouses are not directly operated by Verizon, and given the years of mutual collaboration between the parties, SFS as a project does not consider increasing the DC available capacity to support the new ecosystem.

As Section 3.4 explains, shifting a portion of the DTC demand to the stores will decrease the number of orders arriving at the DTC segment in the DCs. However, if we consider implementing the Optimization Models the arrival rate with which orders are sent to the DC will change as well.

In the Optimization Models, we hold orders for a period of time before making the fulfillment decision and sending them to the stores/DCs. This will create a discontinuous arrival of orders in peaks by period (Figure 15 adds more detail to this concept). For example, let's consider the scenario of 2-hr batching close to the cutoff time (11pm). The orders that arrive between 8-10pm are held for this two hours and at 10pm, when we run the algorithm, several orders will be sent all together to the DC. In one hour, between 10-11pm the DC must handle the load to meet the cutoff time of 11pm. The DC would usually handle these orders over a two to three hour period. Hence, with the SFS and the Optimization Model, the DC could get a surge of orders late in the day, which it might not be able to completely process prior to the cutoff time. The analysis in this section tries to address this concerns and understand the level of disruption SFS brings into the DC operations.

As a baseline assumption, we accept the DC throughput capacity as given by historical data. Once this is obtained, the order arrival rate will be compared to the capacity, so the utilization can be calculated.

5.1.1 Formulation & Parameters

To run a utilization analysis, first we need to understand the capacity constraints. With the warehouse operated by third-party logistics company, historical data was used to assess what is the maximum capacity by hour at the facilities.

We set the nominal capacity equal to the 70th percentile of the maximum orders shipped by hour using historical data of a three-month period. This represents a conservative assumption to estimate the capacity, as it constitutes a volume that the DC can handle. In a three-month period, we know that the DC was able to ship out that level of order volume without needing extraordinary efforts.

For this analysis data from DTC orders only was considered. The DC has a specific segment dedicated to process DTC orders, independently of store replenishments orders. We account for the DTC segment at the DCs as an independent line with its own capacity constraints.

The hourly order arrival rate is calculated as the average of actual hourly data in the same three-month period. Figure 14 compares an aggregate of all the DCs current maximum capacity to the average total order arrival rate, hour by hour over a twenty-four hour day.

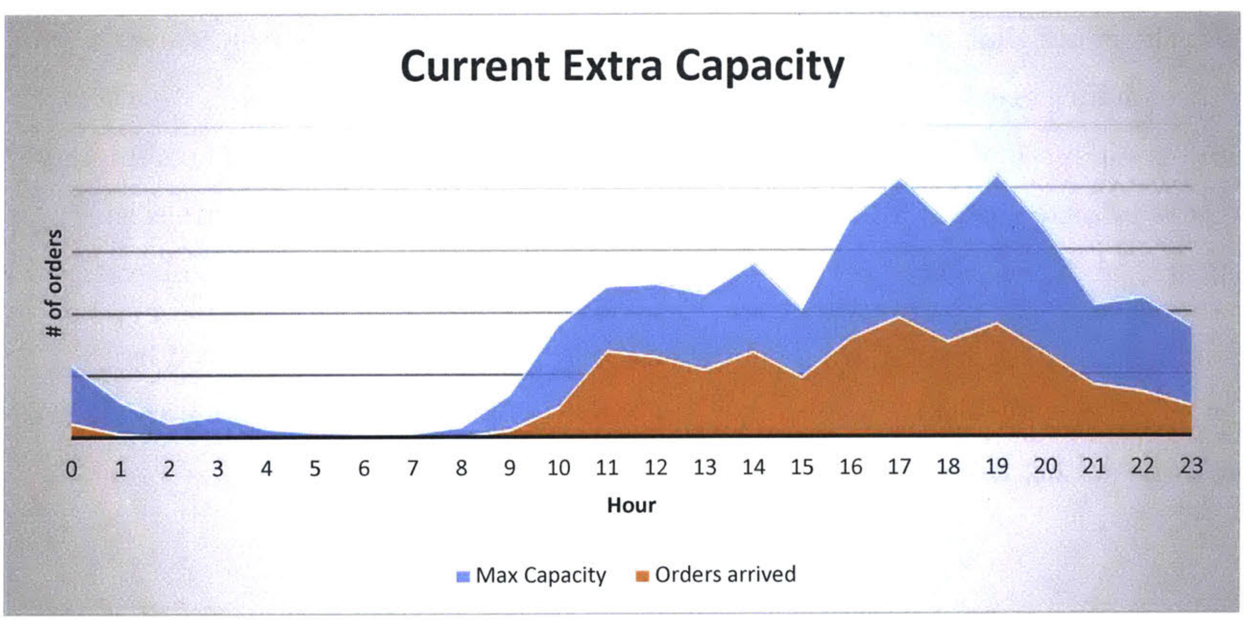

Figure 14 – DC Capacity Analysis

Considering the impact of each one of the routing models, the graph in Figure 15 will be affected in the pattern in which the orders arrive. However, the maximum capacity is considered as given. The Heuristic Model routes orders chronologically as they come and simultaneously with the purchase. While it may have an impact in the number of orders that arrive at the DC, it does not increase the work load or utilization at the DC, as less orders are arriving to the DC (more orders shifted to SFS). With the Heuristic Model, the arrival rate pattern does not change, as the flow of orders is continuous and not discrete.

On the other hand, the Optimization Models hold orders for a period of time to route them together as a batch at the end of each period. This means that orders are not released to the stores or DCs continuously but intermittently. Figure 15 shows the different patterns of order arrival over a day for the three different batching periods.

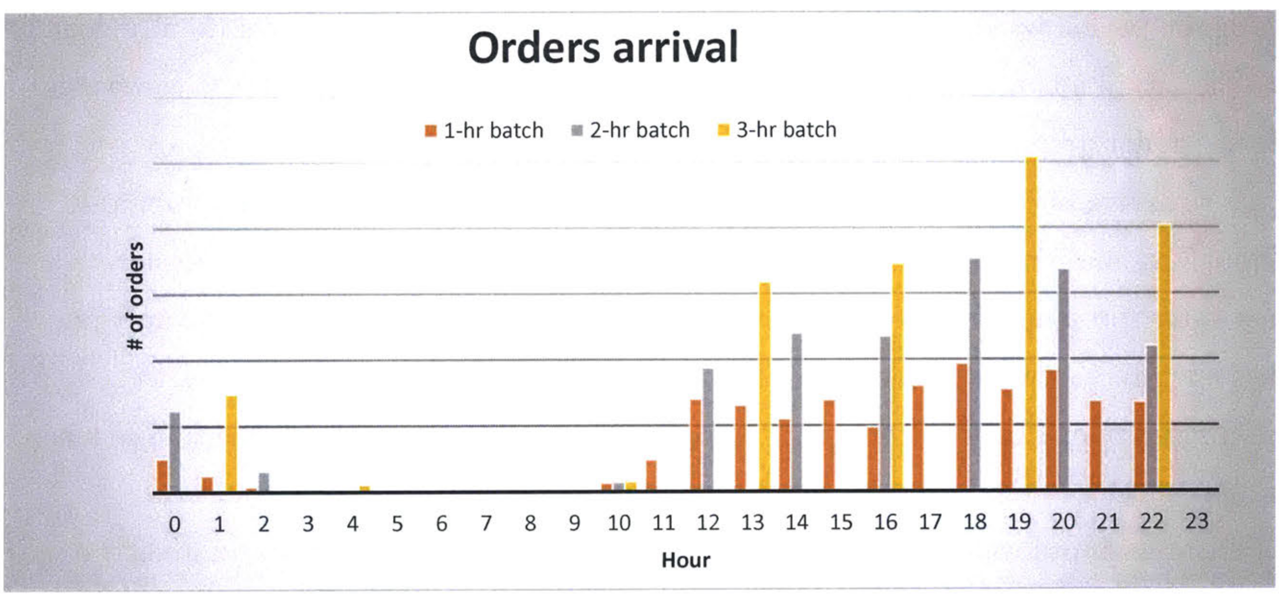

Figure 15 – Orders Arrival Rate

For each scenario, we define an hourly utilization for each hour of the day as the ratio of the average arrival rate (λ) as shown in Figure 15, divided by the hourly capacity given in Figure 14 (see equation 5). When the arrival rate is greater than the capacity we set the utilization to 1 (or 100%), and we assume the arrivals in excess of the capacity flow into the next hour.

$$Utilization = \frac{\lambda}{Max\ Capacity} \quad (5)$$

5.1.2 Results

Using the above specification for the utilization, we can obtain Figure 16 showing the utilization as a function of the hour. In order to compute the utilization, we compare the number of orders arrived in a period to the capacity of the line. For the case of continuous arrivals and one hour batch this computation is straight forward dividing the orders arrived in one hour by the capacity of the specific hour (see Figure 14).

To compute the utilization in the batching scenarios an aggregate of the capacity over the period was used. In this case, orders are held for a period and then released to the DC in a batch. This will shift what capacity we consider in order to compare to the order arrivals. For instance, in the 3-hr batching scenario we compare the orders arrived between 1pm and 4pm to the capacity at the DC between 4pm and 7pm. The challenge arises closer to the cutoff time (11pm). The orders that arrive between 7pm and 10pm are sent to the DC at 10pm, and the DC will have only one hour to process these orders, so in this case three hours of orders arrival are compared to one hour of capacity. As we can see in Figure 17, the utilization reaches its maximum in the period between 10pm and 11pm, and on average several thousands of orders would not be processed until after 11pm, and would miss the cutoff.

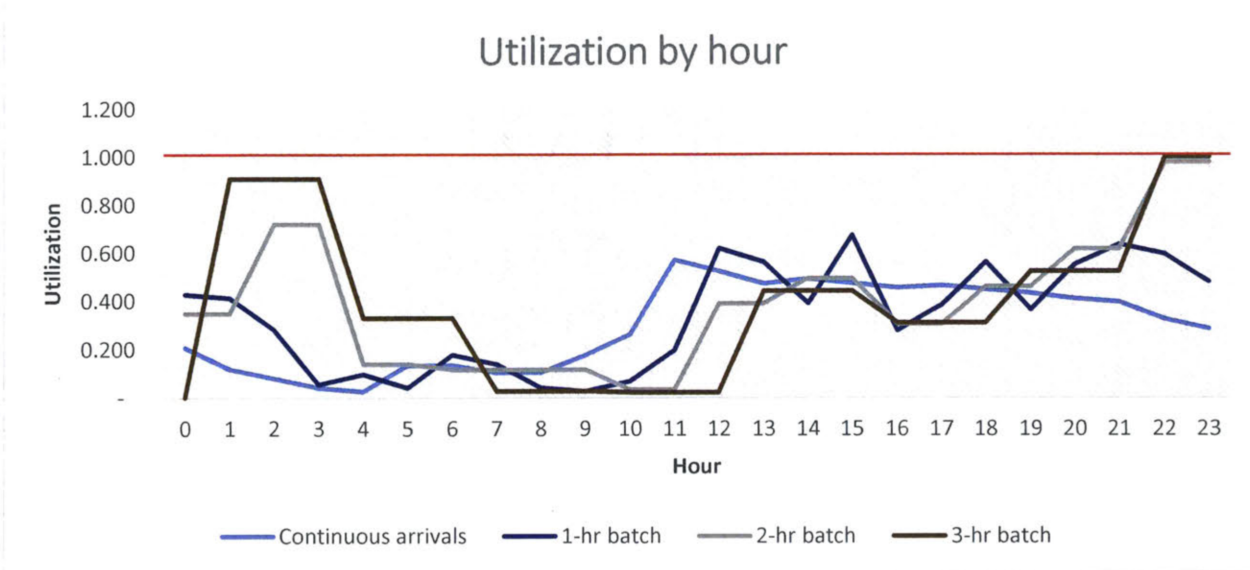

Figure 16 – DC Utilization by routing alternative

The cutoff time constraint is critical, as orders not shipped out by this time will not be delivered on time to the customers and significantly harm Verizon's public perception. Figure 16 proves that as we increase the batching period, the DC capacity starts to be constrained close to the cutoff time, as orders accumulate towards the end of the day.

Although using long batching periods to optimize order routing would be challenging in term of the DC handling close to the cutoff time, we need to weigh in the fact that orders are picked up at stores at 5pm. Thus, stores cannot receive orders after 3-4pm, if we want to make sure the orders can meet the 5pm cutoff and are be properly delivered on time.

A new scenario is analyzed in which, order batching is performed only during stores open hours. After 4pm the orders are automatically routed to the DC per the current model (Section 4.2).

The results presented in Figure 17 compare the maximum utilization across one day between the case of orders batching throughout the day and only within stores open hours. If the Optimization Model is used to route orders only during stores open hours, the DC should be able to handle the modified arrival rate and ship out the orders on time.

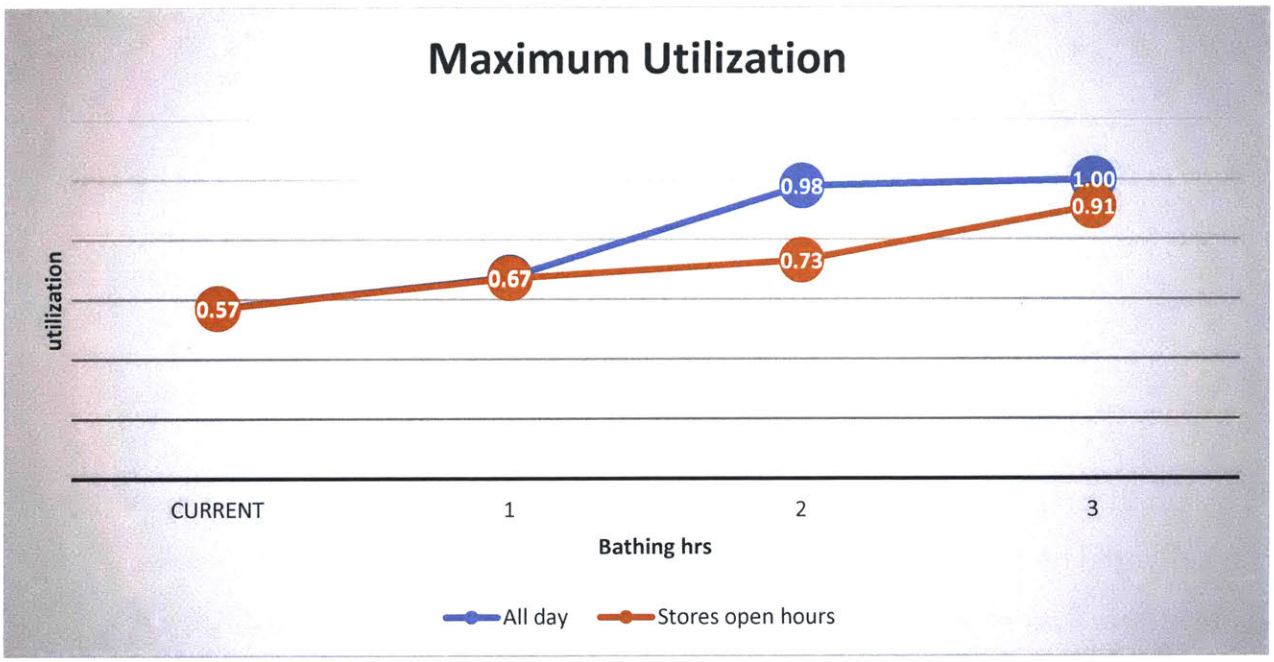

Figure 17 – Maximum Utilization

5.2 DC Safety Stock

Within the company, the Planning team is responsible for determining the inventory targets per SKU across the supply chain. These inventory targets are based on demand forecast, lead times, and review periods. The Planning team estimates the targets basically using the expected cycle stock needed to meet the demand, and the required safety stock to account for demand variability. The demand used to calculate the DCs stock targets is the actual DC throughput. Implementing SFS will certainly decrease at least a portion of the orders directed to the DC, so the expected inventory targets will be affected.

This paper will only focus on analyzing the opportunity to decrease inventory levels in terms of safety stock. The reason for not focusing on cycle stock targets is that until SFS is fully incorporated Verizon needs to make sure it can satisfy the DTC demand and does not want to compromise its service levels. Safety Stock, as its name dictates, provides a safety net that accounts for the variability of the demand, and there is more flexibility to modify it.

5.2.1 Formulation

First, consider the following parameters:

$\Gamma - review\ period$

$L - lead\ time$

$Error_{avg} - \ Average\ error\ (sales - forecast)$

$SL - service\ level$

$\sigma - weekly\ standard\ deviation\ of\ demand$

Using this notation, the Safety Stock target is calculated per equation 6:

$$\boldsymbol{SS_{units} = (\Gamma + L) * Error_{avg} + invnorm(SL) * \sigma * \sqrt{\Gamma + L}} \quad (6)$$

The forecast error can be either positive or negative and this will determine if we are under or over forecasting; we adjust the safety stock accordingly. In the case of negative error the first term of equation 6 will reduce the total amount of safety stock units we hold.

In Equation 6 the target number of units in safety stock for a specific SKU is given by the sum of the average error multiplied by the time it takes the supply chain to react (review period plus lead time), and the inverse normal distribution of the required service level multiplied by the weekly demand standard deviation and the square root of sum of the lead time and the review period.

To estimate how the safety stock changes with an overall decrease in demand at the distribution centers, we need certain assumptions. Based on the simulations performed to obtain the economic benefits of SFS in Section 4, we can estimate the portion of DTC demand that is fulfilled by the stores, which directly decreases the demand experienced by the DCs.

The total demand (throughput) experienced by the DC is composed mainly of DTC demand and Retail store replenishments. Based on historical data we can estimate the percentage of DC throughput that is pertinent to the DTC channel.

The parameter γ represents the portion of demand per SKU that occurs through the DTC channel. The parameter ρ is the portion of the DTC demand that is fulfilled by the stores and not by the DC. The parameter ρ will variate according to how many orders we allow the stores to process per day. As Section 4.8.2 explains, the aggregate number of orders shipped by the stores depends strongly in the constraint of number of orders allowed per store per day.

$\gamma - portion\ of\ DC\ output\ through\ DTC\ channel$

$\rho - portion\ of\ DTC\ demand\ fulfilled\ through\ SFS$

Given the parameters ρ and γ, the decrease in demand experienced by the DCs relevant to the DTC channel and caused by SFS implementation can be calculated.

$$DC\_DTC\_Output\_Forecast = \gamma * Total\_DC\_Output\_Forecast \quad (7)$$

$$DC\_DTC\_Output\_Forecast\_new \; = (1-\rho) * DC\_DTC\_Output\_Forecast \quad (8)$$

For the purpose of this analysis, the percentage of demand through the DTC channel at the DC (γ in equation 7) can be considered constant and is estimated by historical data. Looking at the historical averages of demand through the different channels, γ is set at 15%.

As an example to understand the meaning of the parameter γ, if the total forecast for a specific SKU over a week is 10,000 units from which 3,000 are through DTC and 7,000 to stores, then γ would be 30% (i.e. the portion of the total volume at the DC that is processed by the DTC segment).

Equation 8 estimates the reduction in throughput the DC will experience if SFS in enabled. This reduction is crucial for the calculation, as it will represent the total decrease in demand, and using the new demand forecast we will estimate the safety stock accordingly. The parameter $DC\_DTC\_Output\_Forecast\_new$ is used as the demand forecast to calculate the new safety stock targets to later compare to the original levels.

After calculating the overall demand reduction, we need to make some assumptions in terms of the behavior of the demand distribution after SFS is implemented and a portion of the DTC demand is shifted from the DCs to the stores.

$$\frac{Error_{current}}{Forecast_{current}} = \frac{Error_{SFS}}{Forecast_{SFS}} \quad (9)$$

$$\frac{Var_{current}}{Forecast_{current}} = \frac{Var_{SFS}}{Forecast_{SFS}} \quad (10)$$

Equation 9 assumes that the relative error to forecast ratio will remain constant after SFS is implemented. As we have data for the average error of the current demand forecast, we make this assumption to estimate the error of a portion of the total demand. In other words, if currently the average error of the forecast is 20%, it remains 20% once the forecast and demand decrease by ρ. Having the new demand forecast ($DC\_DTC\_Output\_Forecast\_new$), we can calculate the relevant $Error_{avg}$ to be used in Equation 6.

Equation 10 assumes that the variance of the demand is constant as a percentage of the demand forecast. The variance will then be used to calculate the standard deviation, that also will be used in Equation 6.

5.2.2 Results

Using the results obtained in the simulations from Section 4, and considering the case of a maximum of five orders shipped per store per day, the portion of DTC demand shifted to the stores is $\rho = 9\%$ (equation 8). With this percentage of demand fulfilled by the stores, and using the formulation in Section 5.2.1 the DC safety stock can be reduced by 2.7% in the safety stock targets.

A sensitivity analysis was performed to understand how the change in parameter ρ affects the total WC reduction. For a scenario where 120 orders is the maximum cap for orders per day per store, 30% of the DTC demand is fulfilled by SFS. Figure 18 presents how the safety stock potential reduction changes with the portion of DTC demand shifted to the stores.

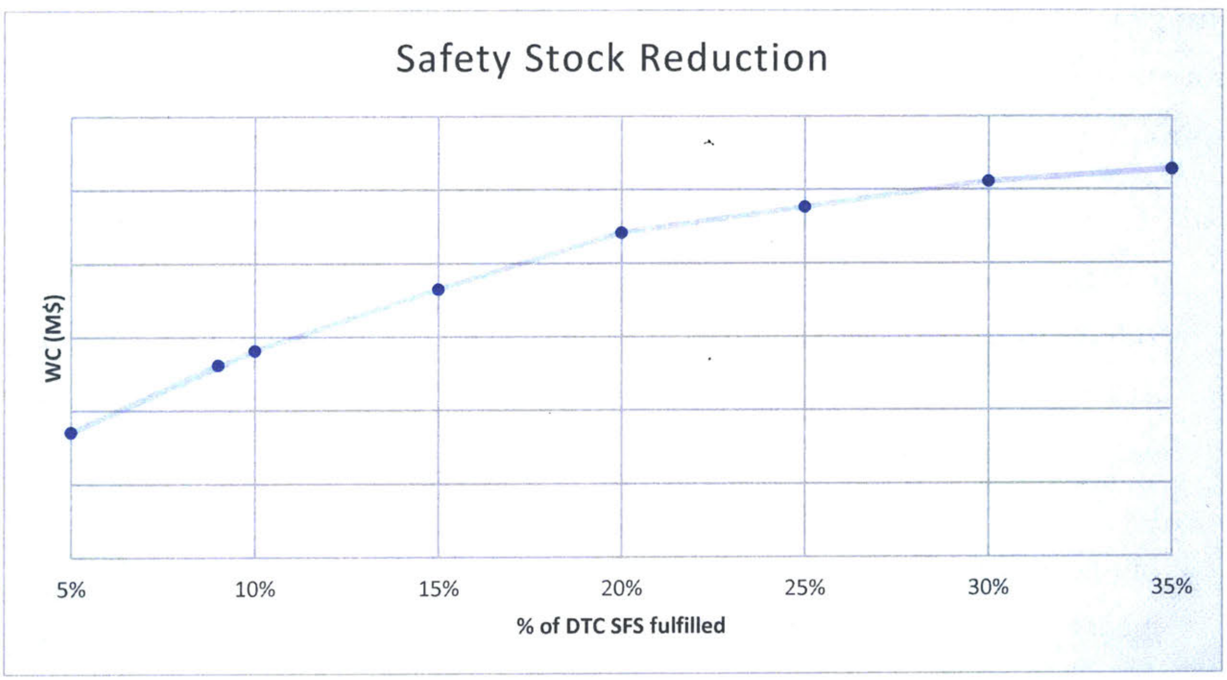

Figure 18 – Safety Stock Reduction

6 Conclusions and Recommendations

This section summarizes the takeaways from collaborative research done during the six-month internship with Verizon Wireless, and provides key recommendations for a gradual and organic implementation of a new routing model.

It starts by synthetizing the conclusions obtained from the mathematical approach explained in this paper, followed by the strategical and tactical implications of a full-scale implementation, and finishes with final recommendations and future steps recommendations.

6.1 Operational Takeaways

Based on the results in Section 4 and 5, there are several key areas where Verizon's supply chain can be optimized yielding major benefits for the company, while enhancing customer experience. The focus of this paper is providing an order-routing model solution, supported by an extensive analytical approach that simulates the behavior of the supply chain. The overall benefits and learnings of each routing model analyzed are:

a) Heuristic Model: relatively simple algorithm based on business rules that routes orders chronologically as they come, based on a greedy algorithm subject to a service constraint at each store. Through the simulation we were able to prove that the Heuristics provides 60% of the improvement in reduced shipping costs from the best theoretical solution (weekly batch) and 84% of the best feasible solution (3-hr batch). This is one of the key findings, as implementing the Heuristic can be much simpler than the more complex models.
b) Optimization Model I: linear program formulation that gathers orders across a time period and routes them together optimizing costs and service levels. This approach obtains an extra 10-15% benefits compared to the Heuristic Model, so it is crucial to understand the added complexity of implementing this model before jumping into final recommendations.

c) Optimization Model II: a non-feasible optimization model that was built to compare the previous alternatives to such an "ideal" solution. As expected, this approach could yield greater benefit as it assumes full demand visibility. Full visibility could, in theory, provide a 15-30% incremental benefit in fulfillment cost reduction; this finding supports the idea that the feasible models constitute a realistically satisfactory approximation of a fully optimized, ideal, supply chain.

6.2 Tactical and Strategic Implications

We need to decide which routing model between the Heuristics and the Optimization Model I should be implemented, and in the case of the latter what batching period should be adopted. This section looks at the different challenges on the way to full implementation.

6.2.1 ATP Check (Tactical)

As Section 5 explains, an ATP check is the test performed by the IT systems to validate real-time inventory levels, once an item is requested. One ATP check is considered when the inventory level of a store-SKU combination is confirmed. The main challenge with the check test is that it may take 3 seconds per test.

Consider the Heuristic Model, where the routing decision is made as the order is being processed (i.e., while the customer is still waiting for order confirmation). The Heuristic Model requires data on real-time inventory for each store-SKU combination to pick a store or DC as responsible for fulfilling the order, and ATP will be the test performed. With the development of e-commerce, customers expect to get order confirmation within seconds. Thus, running this test across each store-SKU combination while the customer is waiting for the order to be confirmed is not ideal.

It is important to comprehend how this is different from the current situation at the company interface. Currently, orders are shipped from either a primary or a secondary DC, so when an item is requested one ATP check (few seconds) is performed to confirm the inventory is there. With the Heuristic Model, we are trying to find the best solution across several hundreds of stores and multiple DCs, so knowing exactly the inventory available at each location becomes challenging.

The way to overcome this kind of challenge is to create a parallel database that keeps track of the real inventory per SKU at each location, and feeds it as an input to the routing model. While, this is not an orthodox way to use the current inventory system at the company, it can be implemented by IT organizations. While several alternatives and approaches were considered when working with IT, the technical specifics of how this is incorporated into the IT systems is not discussed in this paper.

6.2.2 Batching Period (Tactical)

If the Optimization Model I is selected to route DTC orders, we need to determine the period for which orders will be optimized. Section 4.4 explains the difference between the batching periods and Section 4.8.1 shows the final results and benefit for each of them.

The major implication of selecting a batching period is the operational disruption to the distribution centers. While stores will be handling several orders, the DCs will still be responsible for fulfilling the vast majority of orders. Thus, it is crucial to ensure the minimal disruption at the DC level.

Section 5.2.1 explains how the order arrival rates depend on the batching period and provides strong analytical support to guarantee smooth transition to any of the scenarios if we consider optimizing only during store open hours.

Combining the results in 4.8.1 with the analysis in Section 5.2.2, an order batching period of two hours seems the best outcome. It obtains more than 99% of the maximum potential feasible benefits (compared to 3-hr batching), while providing the DCs with more flexibility for order handling throughout the day.

6.2.3 IT implementation (Strategic)

Implementing any of the routing models will require some IT involvement and investment. However, we need to analyze strategically if the expansion to the Optimization Model I is reasonable from the business perspective. In other words, if the marginal benefits obtained by the Optimization Model justifies the additional investment in IT development.

The previous implication, while there are challenges to overcome, tactical solutions are presented on how to resolve them. In this case, it is much more of a strategic decision of whether IT systems has the capability to develop such a model.

The author of this paper does not have access to either the prioritization at the IT organization, or to the actual costs of performing such an expansion. However, this paper states clearly the marginal benefits of such an expansion, so leadership at the company can make a data-driven decision regarding this matter.

6.3 Recommendations and Future Steps

Throughout this paper we analyze the problem statement, provide literature review, explain the methodology and formulations, and analyze the results and implication. Gathering all these together the main recommendations are:

a) Implement the single-item Heuristic Model immediately after SFS is fully implemented. The benefit of restricting SFS to single-item orders is 80-85% of the potential benefit but it simplifies the transition considerably.
b) Expand to multi-item Heuristic Model once the transition is completed and stores have accommodated to deal with orders on a daily basis. This requires additional complexity, so separating this from the initial launch challenge is a good idea.

c) Consider increasing the maximum number of orders per store per day to 20. The benefits can potentially be doubled obtaining operational yearly savings of 9.1% of the total fulfillment costs. Although this is not currently considered as an option, the author recommends to reconsider it in a medium-term once stores are shipping orders consistently.
d) Recalculate Safety Stock levels per SKU at the DC level. Shifting a portion of the fulfillment to the stores yields at least a WC reduction of over 2.7%.

Implementing these recommendations is considered the natural path for Verizon Wireless once SFS is implemented. However, there are some future opportunities that should be considered, to even increase the benefits of this research:

a) Consider the implementation of the Optimization Model I, with a batching period of two hours. Assess the required investment at the IT organization, and compare it to additional potential savings of 0.7% yearly.
b) A similar analysis to the one in Section 5.2 can be performed for cycle stock targets as well. While at this point it is not recommended to decrease DC cycle stock level until SFS is up and running, in the future a major WC reduction can be obtained using these tools.
c) The analysis in this paper was performed using 2016 data. Considering that the DTC channel is growing across almost every industry, we can assume that the benefits will increase accordingly.

7 References

[1] USA Today "Retail stores become shipping hubs to battle Amazon"
http://www.usatoday.com/story/tech/2013/09/28/retailers-ship-from-store/2862405/

[2] Verizon Annual Report 2016 - Available online at:
http://www.verizon.com/about/sites/default/files/annual_reports/2016/downloads/Verizon-AnnualReport2016.pdf

[3] How Verizon, AT&T, T-Mobile, Sprint and more stacked up in Q3 2016: The top 7 carriers - http://www.fiercewireless.com/wireless/how-verizon-at-t-t-mobile-sprint-and-more-stacked-up-q3-2016-top-7-carriers

[4] NRF 2016 Annual Report (2016) – Available online at:
https://nrf.com/who-we-are/annual-reports

[5] Omni-channel retail – A Deloitte point of view (2015)
https://www2.deloitte.com/content/dam/Deloitte/se/Documents/technology/Omni-channel-2015.pdf

[6] Fortna - Ship-from-store distribution: Reducing the impact of an imperfect forecast
http://www.fortna.com/whitepapers/ship-from-store-distribution-reduces-impact-of-imperfect-forecast-FINAL-EN.pdf

[7] Enspire Commerce - "What every omni-channel retailer should know about developing a ship from store strategy"

http://info.enspirecommerce.com/hubfs/enc-wp-ship-from-store-strategy.pdf

[8] Acimovic, Jason, and Stephen C. Graves. "Making Better Fulfillment Decisions on the Fly in an Online Retail Environment." Manufacturing & Service Operations Management 17.1 (2014): 34-51.

[9] Xu, Ping Josephine, Russell Allgor, and Stephen C. Graves. "Benefits of Reevaluating Real-Time Order Fulfillment Decisions." Manufacturing & Service Operations Management 11.2 (2009): 340-355