THE PENNSYLVANIA STATE UNIVERSITY
SCHREYER HONORS COLLEGE

DEPARTMENT OF INDUSTRIAL
AND MANUFACTURING ENGINEERING

USE OF INVENTORY MANAGEMENT FOR LARGE GENERAL ITEM RETAILERS

PATRICK THOMAS LUCAS
FALL 2013

A thesis
submitted in partial fulfillment
of the requirements
for a baccalaureate degree
in Industrial Engineering
with honors in Industrial Engineering

Reviewed and approved* by the following:

A. Ravi Ravindran
Professor of Industrial Engineering
Thesis Supervisor

Paul Griffin
Peter and Angela Dal Pezzo Department Head Chair
Honors Adviser

* Signatures are on file in the Schreyer Honors College.

ABSTRACT

Controlling operating costs is critical in the current competitive landscape of national general item retail chains. Though there are many expenses that fall into the operating budget, inventory control is a major portion of this expense at many large retailers. Having an efficient inventory management policy can provide a strong foundation as many retailers transition to point of sale (POS), radio frequency identification (RFID) and other new technologies to track inventory.

Using recent sales data from a major national retailer, this thesis will apply inventory management theory to construct and evaluate an inventory model for this retail setting. An ABC classification system is used to categorize each Stock Keeping Unit (SKU) based on unit holding cost. The optimal policy is determined to be a base stock system with normally distributed demand. A cost model that calculates the inventory holding cost and shortage cost is optimized to find the optimal order-up-to quantity for each item. This policy is then evaluated through a sensitivity analysis to help understand the major decision factors in the model. The findings indicate that the model will effectively determine the optimal order-up-to quantity and service level for a base stock inventory policy based on the estimated holding cost and shortage cost.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF EQUATIONS

ACKNOWLEDGEMENTS

First I would like to thank my advisor Dr. Ravindran for being my supervisor for this thesis. At a time when I had no topic or knowledge of academic writing he was willing to take me on. Without him I would have been lost throughout the research and writing process and he was very willing to work with me as the timeline tightened towards the submission date.

I would also like to thank Nathan Steadman at Accenture, as well as Lisa and Tom who were my contacts at the client. They all took time out of their normal schedules to help me settle on a relevant topic and then acquire the data necessary to build the inventory model.

Finally I would like to thank my parents, Anne and Tom. Without them I never would have had the opportunity to study at Penn State. They are also the ones who pushed me to apply to the Schreyer Honors College as an incoming freshmen, and then again as a sophomore when I wasn't accepted the first time around.

Chapter 1

Introduction

1.1 Background and Motivation

The topic for this thesis was developed in part through a relationship with a group of Industrial Engineers who work at Accenture in Pittsburgh. The author worked as an intern with this IE group for a total of seven months starting in August of 2011 and concluding in August 2012. While there, he was assigned to a project at one of their major clients that is a retail store with over 8,000 stores in the United States and over 22,000 different product SKUs available throughout the chain. The work performed for this client was part of an ongoing effort to improve on-shelf product stock. This work included improving the methods through which employees keep product on the shelves, including stocking, rotating items, and maintaining system inventory accuracy. Only a very small portion of the in store activities analyzed related to the efficient management of inventory through the order policy.

The current system for managing inventory utilizes a computerized record of the specific store's inventory. This count can easily go wrong but the store has ways of managing it to maintain correct levels. Managers receive an electronic ordering report with the items the computer recommends based on current levels in the system. The algorithm the computer uses is proprietary and will not be disclosed. The manager can then choose to alter the order quantity for each item on the report. This electronic order occurs weekly at the majority of the stores and orders product for the entire store from one of the regional distribution centers. The specific policy at the stores will be discussed later.

Through the continued relationship with an employee of the client, an appropriate project was discussed that would fit the parameters of this honors thesis, as well as accomplish a goal for their organization. A joint decision was reached to develop an inventory policy that will address the management of the client's inventory in hopes of reducing costs across the chain. The client has provided a data set that will serve as a sample set of items tracked from each one of their stores. This data will help determine an efficient ordering policy for the data set, so the end model can be applied to other items at other stores.

1.2 Problem Definition

This thesis is meant to specifically address how an inventory management model can and should be developed for a given data set. It will also address how an inventory model behaves and what should be considered when applying the model across a large set of items. The appropriate inventory policy selected will minimize the total cost for managing the inventory with the given constraints on the current system. The optimum policy is determined by following a fairly standard process that begins by researching what methods are currently being used to manage inventory. The industry knowledge must be combined with a deep understanding of the limitations and opportunities available at the specific store being modeled. A quantitative analysis of the data will determine the parameters of the chosen model, and the optimal conditions for ordering.

The scope of this project is somewhat limited by the proprietary nature of the client information. As seen above, the client has requested to not be identified by name in this paper, and simply referred to as "the client" or "a large national retailer." The process through which the data set was obtained included requesting permission from a managing vice president.

Certain information requested was not made available for the purpose of this thesis and therefore cannot be used to build the model.

1.3 Thesis Outline

This thesis will continue with a literature review in Chapter 2. The literature review will explain why developing an inventory management policy can be beneficial for a business. Since the client is a general item retailer, it will also cover specific aspects about why inventory management is so important in the retail sector. There will also be an explanation of how to quantitatively describe the sales demand and prioritize the classes of items. Lastly, common inventory policies used in industry will be examined, explaining the benefits and drawbacks of each policy.

Chapter 3 will go on to explain the information specific to the model for the client. First, there will be an examination of the type of data obtained from the client. It will call out any drawbacks to the data and identify the direction in which the modeling process will continue. After the data is examined, a distribution will be fit to model the demand for each product, including a calculation of the distributions parameters. Using this information, one of the inventory policies described in Chapter 2 will be applied to calculate the optimal ordering policy that minimizes the total cost of maintaining the inventory.

Chapter 4 will analyze the results obtained in Chapter 3. This will begin by discussing the possible reasons why the final result minimized the total cost. This discussion will identify certain assumptions or parameters from the model that significantly contributed to the end result. These parameters will be further analyzed using a sensitivity analysis. The sensitivity analysis will identify which parameters must be considered when applying the model to other items outside of the ones given in the sample data.

Finally, Chapter 5 will conclude by stating the final recommendations for the client. It will summarize the results obtained in Chapter 3 and call out parts of the analysis from Chapter 4. Using all of this information, a final recommendation will be made as to how the client should proceed applying the inventory policy to the rest of the items in the store.

Chapter 2

Literature Review

2.1 Why Inventory Management?

Every product based company has a need to manage the inventory of its product. Even non-product based companies like accounting firms need to manage the inventory of their product, which would be its employees. The ability to meet consumer demand with a product or service is what separates a successful company from a failing one. At its very core, that is what inventory management is, the ability to balance the delicate relationship between available supply and the customer demand. When the supply exceeds the demand the firm will lose money from holding onto products or employees that aren't being sold or used. When the demand for a product or service exceeds the supply there is a loss of revenue from not meeting that demand. In addition, there is a potential loss of future revenue if the customer need can be met from a competitor. There are costs associated with both of these situations which make it important to find a good balance between the two. We will look further at how to minimize this cost, but first we must understand the types of products we are dealing with.

2.2 Inventory Management for Retailers

The methods of inventory management are different for different industries. For example, in a manufacturing setting the firm must balance raw materials, work in progress, as well as finished goods. This requires an inventory management system that can account for the three product variations. Each of these three categories are at a different stage in the product lifecycle, and thus have different demand requirements within the manufacturing system. These demand requirements largely determine how the inventory system should operate. For the purpose of a retail setting, all the products follow a very similar demand structure since they are all at the final stage of their product lifecycle. Customers enter the store and purchase these items off of the shelf at a seemingly random interval. The specific requirements for modeling this demand will be discussed in detail later, but for now knowing that the customer controls the demand in a retail setting is enough.

Since retailers must interact directly with the customer, they must be located near where the customers live and work. This location constraint creates a significant tradeoff between proximity to the customer and the amount of available space. Densely populated neighborhoods will have a larger demand for retail products, however the cost of leasing space in a densely populated area is significantly more than the cost of leasing space in a very unpopulated area. In a densely populated neighborhood a retail store is going to have limited shelf space and limited backroom space for excess inventory. This space limitation reinforces the need to properly manage this space with an adequate supply of products that customers demand.

General item retailers exist in a very competitive marketplace where price and convenience are imperative to success. Unlike a specialized retailer, general retailers rely on high sales volumes with low profit margin items to stay in business. According to Plunkett Research Ltd (2012), United States retail sales were estimated at $4.920 trillion in 2012 which is roughly a quarter of the 2012 GDP. The sales revenue makes up the vast majority of the total revenue reported for major retailers. Table 2.1 shows the top five earning retailers from 2012 with their reported revenues, incomes, and profit margin. The average profit margin of the top five retailers in the United States is a mere 2.895% and all five are within three percentage points of each other. This industry competition puts immense pressure on retailers to lower prices. The profit margin on the income statement is slightly lower than the actual profit margin of the individual goods sold, but it still calls attention to the fact that retailers are already operating at extremely low product margins. The low product margins mean retailers cannot lower prices any further without reducing the cost of the item itself or lowering the operating expenses of selling that product. Retailers have minimal control over the actual cost of the goods, so they must look for ways to manage operating costs in order to lower prices. Inventory management is a great way for them to reduce operating costs and consequently lower prices.

Table 2.1 Top Five Retailers from 2012. (Data retrieved from Yahoo! Finance).

Company	2012 Revenue	2012 Income	Profit Margin
Wal-Mart	$ 446,950,000,000	$ 15,699,000,000	3.512%
CVS	$ 107,100,000,000	$ 3,461,000,000	3.232%
Costco	$ 99,137,000,000	$ 1,709,000,000	1.724%
Kroger	$ 90,374,000,000	$ 602,000,000	0.666%
Walgreen	$ 71,633,000,000	$ 2,127,000,000	2.969%
Total	**$ 815,194,000,000**	**$ 23,598,000,000**	**2.895%**

The 2008 recession has also led to an increased emphasis on driving costs down. The Mondaq Business Briefing (2010) published a report from the Deloitte Consumer Business Group that discusses the effects of the recession on the retail market. The report showed that, although the adjusted retail sales rose in 2008, many retailers used heavy promotions to maintain sales volume, causing the net profit margin for two thirds of the retailers to decline in 2008. These price slashing efforts and falling profit margins show the growing importance of managing inventory costs in the retail sector as the global economy continues to rebound from the recession.

2.3 Principles of Inventory Management

Before constructing an appropriate model to manage the inventory at a retailer, or anywhere for that matter, it is important to understand what goes into the process of building such a model. The following sections describe the steps relevant for creating an inventory management policy for a general item retailer.

2.3.1 Describe the Demand and Cost Structure

According to Ravindran & Warsing (2013) the first two steps in building an inventory model, are to describe the demand and assess relevant costs. As discussed above, consumer preferences dictate the demand for all retailers so we must try to understand the consumer buying patterns. In a retail store each store reorders inventory separately, so demand must be looked at the individual store level, as opposed to the regional or national level. Using past sales data, (measured daily, weekly, etc. depending on sales volume) the demand patterns can be fit with a known distribution. This distribution will describe the expected buying patterns within a certain range of accuracy.

The costs of maintaining inventory are often described in three parts. The first is the opportunity cost of holding inventory. Inventory holding cost, *h,* is a combined measure of all the costs of keeping excess inventory, including the capital invested in inventory and the rent payments made on storage space. Berling (2005) discusses several shortcut methods used to estimate the holding cost in industry and proposes a more detailed calculation that improves estimation accuracy for the holding cost. The holding cost is traditionally measured in dollars per unit per time.

The second relevant cost is the cost associated with placing an order. Order cost, *A*, includes everything associated with placing and receiving an individual order. If the shipping costs are a flat rate for an entire order, the order cost is often modeled as a constant. If the order cost varies by size, modeling the ordering cost becomes more complex because it is now a function of the order quantity.

The third cost is the out of stock cost that measures the cost of sales lost due to stock-outs. It can be quantified as the profit lost from each item that was unable to fulfill the customer demand. This is the most difficult to measure because customers won't always inform the store they couldn't find the item they intended to buy. Some customers know exactly what they want, and if they don't find it they will go to another store looking for it. Other customers will purchase items on impulse when they see them at the store. If the impulse product is out of stock, the store and the customer will never know that the demand for that product ever existed. A common practice in industry is to set a safety stock quantity per item based on desired service level. Setting a high service level protects against a stock out and minimizes the profit lost due to unfulfilled demand. Service level, whether calculated or determined by the corporate strategy, is a measure that must be set for all items.

2.3.2 ABC Classification

The buying patterns will vary greatly by item. In retail settings with thousands of SKUs it is neither possible nor cost effective to develop an inventory model for each individual SKU. A common industry practice is to divide the SKUs into three categories based on sales volume. This classification method is referred to as an ABC analysis. The most common practice is to have class A be the top 20% of items comprising of 80% of the sales dollars. This 80-20 relationship is described as the Pareto Principle (Ravindran & Warsing, 2013). It was originally stated by

Italian economist Vilfredo Pareto to describe the distribution of income in society (Femia & Marshall, 2012). When applied to retail sales, this principle estimates that the top selling 20% of products account for 80% of the revenue in the store. Class A items are deemed the most important and require the most attention of the three. Class B items are the next most important and make up the next 30% of items, typically comprising of 15% of the total revenue. Finally, class C items make up the last 50% of items sold, and account for only about 5% of the sales revenue. The division of the classes of items can be seen in the graph of cumulative percent of sales and cumulative percent of items shown in Figure 2.1. The vertical lines have been added separate each class of items with class A on the far left, class B in the center, and class C on the right.

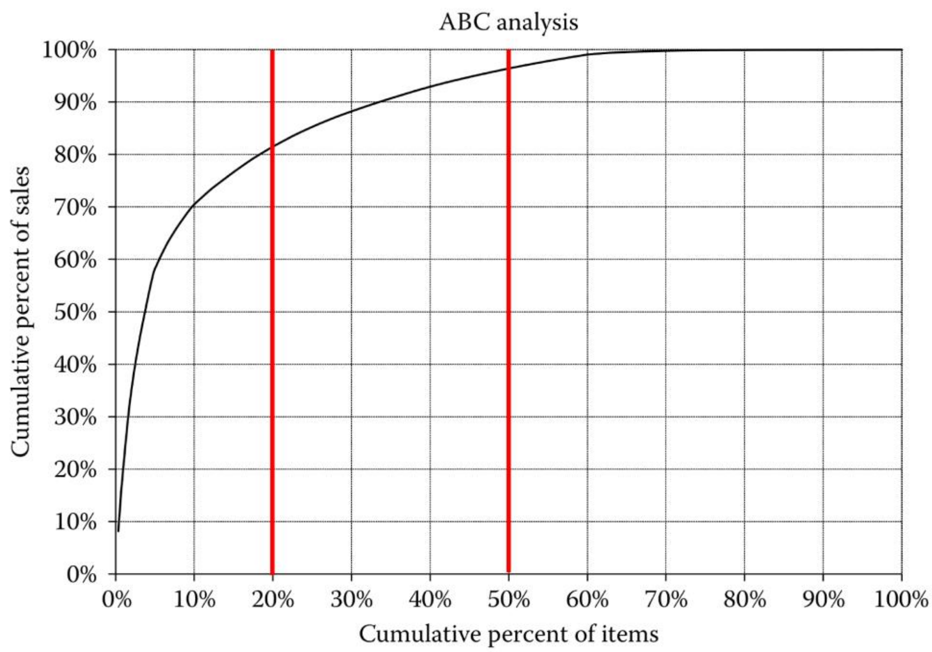

Figure 2.1 ABC Inventory Analysis Plot (Ravindran & Warsing, 2013).

Other classification methods can be implemented that utilize a measure other than sales volume. Almost any measure can be used to classify the data, as long as it matches the priorities of the company. Teunter et al. (2010) suggest that there is an optimal method for establishing a class system based on a measure they call *cost criterion* that minimizes the inventory management costs. However the items are classified, this is an important step in prioritization to make sure more important items receive more attention.

2.3.3 Inventory Policy Selection

The different classes of items are usually managed with one of two different types of inventory policies. The first policy is the continuous review or (Q,R) policy. This policy constantly monitors the inventory level of a product and suggests a reorder of a quantity, *Q*, when the inventory reaches a specified reorder point, *R*. This policy is optimized when the total cost of ordering and holding inventory is minimized. The total annual cost equation for this policy can be expressed as

$$\mathrm{TAC}(Q) = OC + HC = A \cdot \frac{D}{Q} + h \cdot \frac{Q}{2}$$

Equation 2.1 Total Annual Cost of a (Q,R) Policy.

where *OC* is the order cost and *HC* is the holding cost. The order cost depends on, *A*, the cost per order placed, times number of orders placed per year, shown as the yearly demand, *D*, divided by, *Q*, the order quantity. The holding cost depends on, *h*, the holding cost per unit per year times the average inventory level, which for constant demand is given by half the order quantity.

The quantity that minimizes this cost is known as the economic ordering quantity (EOQ). This is derived by taking the partial derivative of the total cost equation (1) and setting it equal to zero. Solving for the optimal order quantity, *Q**, or economic order quantity (EOQ), gives the equation in (2).

$$Q^*_{EOQ} = \sqrt{\frac{2AD}{h}}$$

Equation 2.2 Efficient Order Quantity for a (Q,R) Policy.

These equations are for very basic continuous review policies with constant demand and lead time. When introducing uncertainty, the equations must be modified to account for the variations in lead time demand. The modification depends on how the demand and lead time are modeled.

Continuously monitoring inventory levels was a challenge before point-of-sale (POS) monitoring systems were widely available. These systems maintain a computerized record of the inventory level that update when products are scanned at the register. POS inventory systems aren't perfect however, and can be time consuming to maintain accurate levels. As one might expect, continuous review policies are easier for scenarios when each product is ordered separately and can each be ordered at the critical quantity. When several items are ordered together as part of one shipment, (the case at most large retailers), a continuous review policy is more difficult to implement. The order cost and holding cost terms in the total cost equation are transformed to include all the items being ordered, and the total cost is aggregated across all the items. The problem is then concerned with optimizing when to order and not how much to order. See Ravindran and Warsing (2013) for more detail and the full model for multi-item inventory problems.

The second policy is the periodic review policy. There are three common versions of the periodic review policy. The first is the reorder point, order-up-to policy, also called an (s,S,T) policy. This policy reviews inventory levels at the end of each period, *T*, and if the item count is below the reorder point, *s*, an order is placed to replenish the inventory to the order-up-to point, *S*. In this policy, as in the continuous review policy, the idea is to minimize the total annual cost of managing inventory. This minimum is computed by solving for the optimal parameters s^* and S^*separately. Ravindran & Warsing (2013) describe that the optimal quantities for the reorder point and the order-up-to quantity require sophisticated mathematical computation, but a close approximation can be obtained using calculations similar to those for the continuous review policy. In most periodic review systems, the review period is determined based on supplier constraints and is usually known. The reorder point, *s*, should be set at a level that protects against stock outs during the review period plus the lead time for new products to be delivered. When there is uncertainty in the demand or the lead time, that uncertainty can be accounted for by setting a desired cycle service level (CSL). For normally distributed demand the optimal reorder point is given by

$$s = \mu_{\mathrm{DLTR}} + z_{\mathrm{CSL}} \sigma_{\mathrm{DLTR}}$$

Equation 2.3 Optimal Reorder Quantity for an (s,S,T) Policy.

where μ_{DLTR} is the average demand during the review period, the lead time and z_{CSL} is cycle service level factor, and σ_{DLTR} is the standard deviation of the demand.

The order-up-to point *S* is then determined as the reorder point plus the optimal order quantity. A simple way to determine this order quantity is to use the EOQ calculation from the (Q,R) policy. Therefore the optimal order up to quantity can be expressed as

$$S = s + \sqrt{\frac{2AD}{h}}$$

Equation 2.4 Optimal Order-Up-To Quantity for an (s,S,T) Policy.

This policy is optimal when continuously reviewing inventory levels is a costly activity. For class B or C items the cost of continuously reviewing inventory levels could be significantly more than the cost of switching to a periodic review policy. The periodic review policy also makes it easier to manage several SKUs in a single order because there is a fixed review period for each item. In this scenario the parameters *s*, and *S* would be evaluated for each item and the order would include only the items which need to be ordered.

In another variation of the periodic review the order cost is negligible or zero. In this scenario an order is placed each period, which allows average inventory levels to be lower than in an (s,S,T) policy. For a retailer ordering thousands of different SKUs, there is likely to be an item that requires a reorder in each specified period. Therefore, the cost of ordering an additional item in the same shipment is negligible and the inventory policy can be modeled as having an order cost of zero. When the order cost can be modeled as zero, there is no reorder point, *s*, and it is just referred to as an order-up-to (S,T) system. For an order-up-to system the order quantity for each item is equal to the order-up-to level, *S*, minus the current inventory level, *x*, for that item. Assuming a fixed review period, only the order-up-to quantity, *S*, has to be determined. The optimal order up to quantity follows the same logic as the reorder point equation in (2.3) in that it only has to protect against a stock out during the review period and the lead time. Therefore the optimal *S* for normal demand can also be expressed as

$$S = \mu_{DLTR} + z_{CSL}\sigma_{DLTR}$$

Equation 2.5 Optimal Order-up-To Quantity for an (S,T) Policy.

The last version of the period review system is for the specific case of the (S,T) policy where the review period is equal to the smallest discrete time unit in the system, or in other words T=1. This policy is called the base stock system or a (S,1) Policy. In the base stock system an order is placed each period to replenish the inventory up to the order-up-to quantity. The order-up-to quantity is determined in the same way that it is determined for an (S,T) policy. Because the review period is equal to one, the normally distributed expression for *S* can be rewritten as

$$S = \mu_D(L + 1) + z_{CSL}\sigma_D\sqrt{L + 1}$$

Equation 2.6 Optimal Order-Up-To Quantity for an (S,1) Policy.

where z_{CSL} is, again, the cycle service level factor determined from the normal distribution.

The appropriate selection of an inventory policy largely depends on how inventory levels are maintained and how product is ordered for the particular company. Certain companies have more flexibility with suppliers, and can continuously monitor inventory. This would allow them to become very efficient using a continuous review policy. Other companies are restricted by suppliers and inventory resources, making a periodic review system more economically feasible.

Chapter 3

Model Application and Validation

3.1 Data Description

As stated in the introduction, the data for this model comes from a large national retail chain. There are eleven items tracked in the data set that were selected by an employee of the retailer to represent a variety of items throughout the store. The collection occurred weekly and lasted one year starting on 7/7/12 and ending 7/6/13. The items have been classified as either a basic item, meaning it is available throughout the year, or given a seasonal label. Table 3.1 shows the sales summary for each item across the entire supply chain.

Table 3.1 Item Summary of Sales Across Entire Supply Chain from 7/7/12 – 7/6/13.

Item	Season	Sales Units	Sales Dollars	Unit Profit	Inventory Holding Cost ($/unit/yr)	Lead Time (days)
1	BASIC	1,255,284	$177,099,678	$26.59	$11.84	6
2	BASIC	9,666,209	$6,106,458	$0.53	$0.05	17
3	BASIC	606,793	$90,833,704	$36.01	$11.40	7
4	FALL	22,028	$594,527	$9.94	$2.01	0
5	BASIC	108,915	$425,205	$1.92	$0.21	16
6	BASIC	322,837	$564,474	$0.63	$0.14	11
7	BASIC	21,501	$221,353	$9.03	$0.15	13
8	BASIC	57,425	$285,218	$3.34	$0.15	9
9	BASIC	300,527	$529,389	$0.63	$0.14	11
10	BacktoSchl	236,668	$3,164,949	$4.49	$0.95	9
11	BacktoSchl	83,405	$3,462,599	$9.20	$3.08	9

Since this is only a very small subset of all the items available in the store it is not practical to partition these items by sales volume or dollars because it is not known how much of the whole pie each item represents. Also, in this subset there would not be a fair division amongst the items by trying to categorize them in terms of sales units or sales dollars. Item 1 dominates the sales dollars category on its own with 63% of the total dollars earned, and item 2 dominates the sales volume with 76% of the overall sales. Instead items will be prioritized based on the holding cost. Looking quickly at the data above it is easy to see that there is a large discrepancy in the holding cost between certain items, but the distribution of holding costs in this subset maintains the Pareto principle of an 80-20 relationship. Also, when prioritizing the items by holding cost the store is emphasizing the items that are expensive to keep in inventory, which is one way to help reduce costs. Using the 20-30-50 method for classes A, B, and C respectively the items can be sorted into the categories summarized in Table 3.2 below.

Table 3.2 ABC Analysis by Holding Cost

Class	Items	% of Items	% of Holding Cost
A	1, 3	18.2%	77.2%
B	4, 10, 11	27.3%	20.1%
C	2, 4, 6, 7, 8, 9	54.5%	2.7%
Total		100.0%	100.0%

The national data shown in Table 3.2 is useful for the ABC classification, however it is much more practical to use data from an individual store to help determine an inventory policy. The client has provided data for each of the eleven items from two stores in the chain. The detailed summary of this data can be found in Appendix A, but for the purposes of the model the yearly sales data for each store for items 1 and 3 are shown below in Table 3.3.

Table 3.3 Yearly Sales Quantity for Items 1 & 3 by Store.

Item	Store	
	A	B
1	117	88
3	46	94

Items 1 and 3 are the two items that make up class A, meaning they are prioritized in the classification scheme. When looking at the week by week sales data for each item (available in

Appendix A), these two items also had the most consistent sales patterns throughout the collection period making the demand easier to model. Other items had strong seasonal trends or, (in the case of items 2 and 4-9), were introduced in the stores in the middle of the collection period causing the data to be strongly skewed towards the final six months of collection. The remainder of the analysis in this thesis will focus on these two items in class A for these reasons.

Using the week by week data provided by the client, the demand patterns can be analyzed and modeled at each store. One potential issue with the raw data is that includes returned items by adding a negative sale. If the return occurs in the same period as the original sale then it goes unnoticed, however when the return occurs in a subsequent week, and no other units are purchased during that week the data shows a value of -1. This is problematic because the sales data is by definition discrete, meaning that only whole units of inventory can be purchased. To avoid the problem with returns, all negative numbers in the raw data have been changed to zero. For data with low demand, the first distribution to test for a fit is a Poisson distribution. The Poisson distribution describes discrete events with an exponentially distributed inter-arrival time. One of the properties of this distribution however is that the variance must equal the mean. The weekly mean, variance, and standard deviation for demand are shown in Table 3.4. Looking at the mean and variance for the data, with the exception of item 3 at store A, it is clear that the Poisson distribution will not accurately model the data.

Table 3.4 Weekly Sales Summary (negative sales removed)

Measure	Item 1		Item 3	
	Location A	Location B	Location A	Location B
Weekly Average	2.2264	1.6792	0.9434	1.7736
Variance	3.9110	2.5198	1.1100	3.5714
Std. Deviation	1.9776	1.5874	1.0536	1.8898

When a Poisson distribution fails it is common to check for a normally distributed demand. The normal distribution is described by the mean and standard deviation shown in Table

3.4. The graph in Figure 3.1 shows a histogram for item 1 at store A of observed demand frequency and a normal distribution of mean 2.264 and standard deviation of 1.9776 overlaid on top.

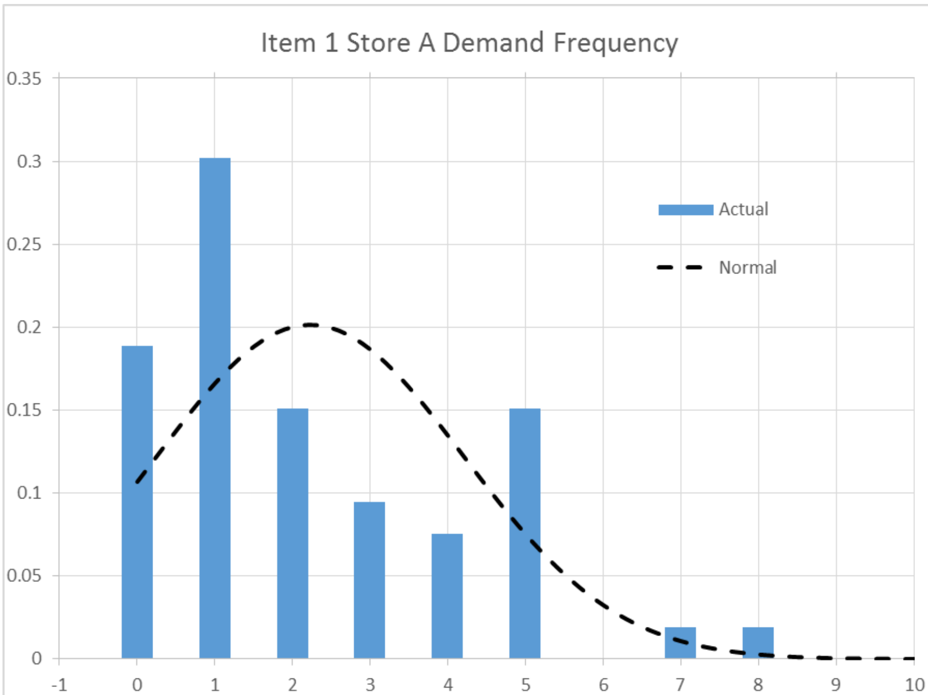

Figure 3.1 Histogram of Item 1 Store A Demand Frequency

Although the normal distribution doesn't overlap perfectly with the observed frequencies, it appears to be a good fit. Similar plots for the other items were created and also appear to fit the normal distribution. The plots for other items are located in Appendix B. As mentioned, the demands for items 1 and 3 are fairly consistent throughout the year, which also is an indicator for normally distributed demand. Another less common distribution may have a more accurate fit, but for simplicity the demand for all items will be assumed normally distributed.

3.2 Policy Selection

The next step is to establish an inventory policy to meet the needs of the store. The stores in this retail chain order items from one of their regional distribution centers. As described in Chapter 2, retailers with thousands of SKUs can simplify the ordering process by using a periodic review process assuming a zero order cost. The sales data is given in weekly increments, and due to the upstream optimization of the supply chain, it is only feasible for each store to order at most once per week (a select few high volume stores order twice per week). Attempting to change the way the entire chain orders may have an unforeseen negative effect on the overall supply chain. The two stores in this data set have an order period of one week, or T=1. As described in Chapter 2, when the order period is the equal to the smallest discrete interval of the sales data, a base stock policy should be implemented.

For a base stock system with a known order period length, only the optimal order-up-to level, *S*, needs to be determined. When there is a set service level, the optimal order level, *S**, is determined using Equation 2.6. When there is no set service level, *S** can be found by minimizing the total cost for maintaining the inventory. The total cost includes the total inventory holding cost and the total cost of lost sales due to shortages.

3.3 Calculations & Results

In order to properly model the cost, it is necessary to first compute the expected weekly demand for the order period. In base stock system the order must account for the product demand during the review period and during the product the lead time. Going back to the information in Table 3.1, it can be seen that this lead time is constant. This simplifies the model because lead time variation won't have to be accounted for. The graph in Figure 3.2 shows the general relationship of the inventory level over time in an order-up-to system.

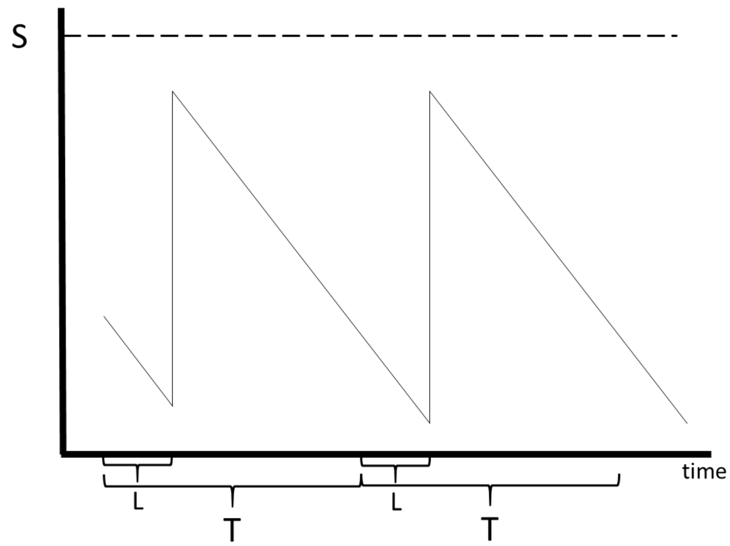

Figure 3.2 (S,T) Inventory Policy.

As shown in the graph, when an order is placed at the end of each period, the inventory level continues fall until the product is delivered after the lead time. Therefore, instead of modeling the average weekly demand for each item, the average demand for period T + L should be used in the model. Modeling for the whole ordering period is what Jensen & Bard refer to as modeling for inventory position instead of inventory level (2003). The inventory position is the level of on hand inventory plus the quantity that has been ordered and not yet delivered. For normally distributed demand, the average demand for this period is μ_{DLTR}, and the standard deviation is σ_{DLTR}. In a base stock system the period T is equal to one, so μ_{DLTR} and σ_{DLTR} can be calculated as

$$\mu_{DLTR} = \mu_D(1 + L)$$

Equation 3.1 Average Demand During Lead Time and Review for (S,1) Policy.

$$\sigma_{DLTR} = \sigma_D \sqrt{1 + L}$$

Equation 3.2 Standard Deviation During Lead Time and Review for (S,1) Policy.

Before calculating these parameters, a decision needs to be made about the lead time for items 1 & 3. Because the items are being ordered together, they must be reviewed at the same time. To be safe the larger of the two lead times is chosen, which is 7 days for item 3 (Note that the item 1 lead time is only one day shorter at 6 days. The effect of the lead time will be examined further in the sensitivity analysis). Since the demand is measured in weekly increments, *L* is set to equal one. Using the original mean and standard deviation from Table 3.4 in Equations 3.1 and 3.2, the resulting calculations for μ_{DLTR} and σ_{DLTR} are shown in Table 3.5.

Table 3.5 Average Demand for the Ordering Period T+L.

Measure	Item 1		Item 3	
	Location A	Location B	Location A	Location B
μ_{DLTR}	4.4528	3.3585	1.8868	3.5472
σ_{DLTR}	2.7968	3.1747	2.1071	3.7796

Using the equations for the base stock system presented in Chapter 2 an optimal order-up-to level can be calculated given a cycle service level (CSL). For example, if the service level for all items is 95%, S^* can be calculated using Equation 2.6. This calculation is shown in Table 3.6. Since the order up to level must be a whole number, the values for S* in Table 3.6 must be rounded up to the next whole number to ensure the service level requirement.

Table 3.6 Optimal Order-Up-To Quantities for 95% CSL.

Item	S*	μ_{DLTR}	$Z_{0.95}$	σ_{DLTR}
Item 1 Location A	9.053	4.4528	1.645	2.7968
Item 1 Location B	8.580	3.3585	1.645	3.1747
Item 3 Location A	5.353	1.8868	1.645	2.1071
Item 3 Location B	9.764	3.5472	1.645	3.7796

As mentioned, the goal is to determine an order level that minimizes the total cost. The calculation above simply selects the optimal S based on a desired service level. However, if an estimate for the cost of lost sales due to stock outs can be calculated, an optimal policy can be developed to minimize the total cost and calculate an optimal service level. The total cost is equal to the sum of the cost of maintaining excess inventory and the cost of lost profit due to shortage. In a given period there is a probability *f(x)* described by the normal probability density function (pdf) that *x* number of units will be demanded by the customers. The combination of all probabilities multiplied by the number of units that are short or in excess will return an average quantity of short and excess estimation for the given time period. For a continuous range of possibilities this combination of all possible demand is represented as the area under a curve. Definite integrals with lower and upper bounds are used to calculate the area under a curve. The average quantity described by the area is then multiplied by the holding cost and the shortage cost to create the total cost equation. This total cost equation as a function of *S* is shown in Equation 3.3.

$$TC(S) = h\int_0^S (S-x)f(x)dx + C_s\int_S^\infty (x-S)f(x)dx$$

Equation 3.3 Total Cost as a Function of Order Quantity

Where:
h = holding cost ($/unit/period)
x = the actual quantity demanded per period
S = the order-up-to quantity
f(x) = the normal probability density function
C_s = shortage cost ($/unit)

The first of the two integration terms determines the expected number of excess units held in inventory after a given period. Similarly the second integration term determines the expected number of units demanded but unfulfilled. Both of these terms are then multiplied by their associated cost rate and then combined to give the total cost. It is important to note how the bounds of these integration terms are set. The inventory holding cost is only relevant when the order level *S* is greater than the actual quantity demanded, *x*. If the actual quantity demanded were greater than the order level a negative dollar amount is calculated and therefore must be ignored. The same situation occurs with the shortage cost for values of *x* less than *S*. When the demand and the order level are exactly equal both costs are zero because there will be no lost profit from missed sales, and there will be no excess inventory. If consumer demand was constant and known the optimum solution would have a zero total cost.

The right integration term has an upper bound of infinity, which may seem alarming, but the limit of the probability function as *x* approaches infinity is zero. Technically there is a probability that an infinite number of customers will show up, however this probability is exactly equal to zero for the normal distribution. As *x* increases the probability term decreases at a significantly faster rate than the *(x-S)* term increases. This allows the integration term to be evaluated and makes it possible to calculate the cost for the infinite number of possible *x* values.

As mentioned, the normal probability density function *f(x)* describes the probability of exactly *x* units demanded during the order period. This probability function is given by the following equation with μ as the mean and σ as the standard deviation for the demand.

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} exp\left\{-\frac{(x-\mu)^2}{2\sigma^2}\right\}$$

Equation 3.4 Probability Density Function for a Normal Distribution

The common mathematical solution for an inventory optimization problem is to take the partial derivative of the cost equation and set it to zero. However when substituting the expression in (3.4) into *f(x)* in (3.3), the integration becomes very complex. It would require some advanced transformations using the error function to integrate. Instead, a program like excel solver can iteratively find the optimal solution. Table 3.7 shows a sample calculation for the total cost of managing item 1 at store A when S is set to 7.

Table 3.7 Excel Calculations for Item 1 at Store A.

Store	**A**	**Item**	**1**
μ_{DLTR}	**4.4528**	σ_{DLTR}	**1.9776**
Order Up to Level, S	**7**		
Excess Cost		**Shortage Cost**	
Average Units Excess	2.3664	Average Units Short/review	0.2753
Unit Holding Cost/yr, *h*	\$11.84	Unit Shortage Cost, C_s	\$26.59
Yearly Excess Cost	\$28.02	Review Periods/year	26.07
		Yearly Shortage Cost	\$190.86
Service Level	81.88%	Total Annual Cost	\$218.88

This result is not the optimal result but is provided for the sample calculations. The information in the header of the table is set for each item at each location. The average demand and standard deviation for the order period is taken from the earlier calculations in Table 3.5. The current order quantity, *S*, is entered manually. Changing this value will automatically update all of the remaining calculations. The table is set up in two halves with the left hand side calculating the excess cost and the right hand side calculating the shortage cost. The first calculation on both sides is to evaluate the definite integrals for the given order quantity. These values represent the average quantities excess or short for a given order period and can be calculated in excel using a trapezoidal area function for area under a curve. This calculation estimates the area by calculating a thin sliver of area under the curve. For a given function *f(x)* the area shown in Figure 3.3 is described by the equation in (3.5).

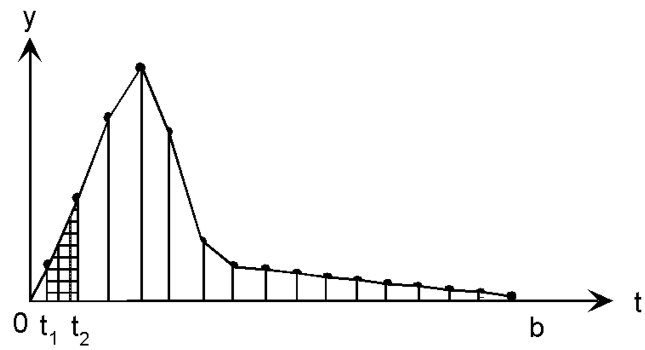

Figure 3.3 Trapezoidal Area Function for *f(x)* (Haggerty, 1999).

$$Area = (t_2 - t_1)\left[\frac{f(t_1) + f(t_2)}{2}\right]$$

Equation 3.5 Trapezoidal Approximation for Area Under a Curve

As the step size t_2-t_1 gets smaller the accuracy of the estimation increases. For the calculations in Table 3.7 a step size of 0.1 was used allowing for negligible error

The average quantity excess for the period will be the constant throughout the year and can be multiplied directly by the yearly holding cost to get the annual cost for excess inventory. The average quantity short for the period must be extended throughout the year. Since there are just over twenty six two-week periods in a year, the average quantity short per period is multiplied by the quantity 26.07 to result in the average quantity short per year. This value is then multiplied by the cost per unit short to get the annual shortage cost. These two values are then added together to calculate the total annual cost of the policy. The service level can also be calculated using the normal distribution function in Microsoft Excel, and measures the cumulative probability that demand will be less than the given value of S.

The Excel solver tool is set up so it can minimize or maximize the value of a particular cell while varying the value of another cell or set of cells. For this model, the total annual cost cell will be selected as the target for minimization. The order quantity cell will be selected as the cell that can be varied. The solver tool also requires the user to declare any constraints on the system. Since the inventory system requires a whole number for the order up to level, the value for *S* must be an integer greater than zero. Applying this optimization software to each item at each location yields the optimal order-up-to values shown in Table 3.8.

Table 3.8 Excel Solver Optimization Calculations.

Measure	Item 1		Item 3	
	Location A	Location B	Location A	Location B
Order Quantity	10	10	7	12
Annual Cost	$75.79	$73.45	$46.08	$84.27
Service Level	97.634%	98.178%	99.238%	98.734%
Safety Stock	5.55	6.64	5.11	8.45

For the optimal result all of the service levels are greater than 97.5% meaning that there is less than a 2.5% chance of missing a sale in a given week, or 1 missed sale for every 40 (or more) weeks. The safety stock is given in partial units, because it is calculated as the expected demand subtracted from the order level. Since the expected demand is an average that can take on partial values, the safety stock can also be viewed be viewed as an average for the year. The measures calculated so far have represented the annual quantities and costs even though the original cost equation calculated the cost per period. In practice, there would not be both a shortage and an excess cost in a given period. Therefore it is not as useful to present the total cost of the policy in terms of one period. Over the course of a year the demand will fluctuate each period, creating both overages and shortages. The cost of each shortage and overage will accumulate, and in the long run, should represent the value given for the total annual cost.

Chapter 4

Model Analysis

4.1 Analysis of Results

The results from the optimization merely give the answer to the question regarding what is the optimum order quantity to reduce total annual cost based on the assumptions in the model. In order to apply this model to other items and other stores it is necessary to analyze deeper as why this policy returned the results found in Chapter 3.

In order to understand the effect that the order quantity has on the total cost of managing inventory, the total cost has been plotted against the order quantity in Figure 4.1.

Figure 4.1 Plot of Total Cost vs. Order Quantity for all items.

When the total cost is plotted against the order quantity, it is clear that in a situation when the order quantity is less than the optimum the total cost is extremely high. The cost for holding excess inventory then slowly increases as the optimum order quantity is surpassed. It should be noted that the true minimums of these curves occur in between the integer values, however since a partial value of a product cannot be ordered, the integer order quantity with the lowest cost is selected.

To better understand how the shape of the total cost curve is developed it is also useful to plot the total cost curve on top of the cost curves for shortages and excess. This graph for item 3 at store A is shown in Figure 4.2. This curve was to represent all four of the item breakdowns because it is closest to the origin and is easier to plot, however all four curves take the same general shape.

Figure 4.2 Cost Breakdown for Item 3 Store A.

For order quantities lower than the optimum of 7, the shortage cost is the main driver for the total cost, with the excess cost representing a very small fraction. When the order quantity is below the optimum, there is a very large probability that there will be a shortage of multiple items. This shortage will cause a loss in profit for every item short. The curve is not as steep as the order quantity surpasses the optimum. There is less of a probability for a shortage of items beyond this point, so almost all the demand is being fulfilled. It continues to rise at a rate consistent with the cost holding one more unit of inventory per year because it is unlikely that this additional product will be sold, and will sit in inventory. At the optimum order quantity of 7, the excess cost accounts for around 90% of the total cost. This clearly shows that the cost of keeping excess inventory is much lower relative to the cost of missing a sale.

The items used were the highest profit items from the original eleven items in the data set. This could possibly lead to a skewed result in favor of keeping excess inventory so as to not lose sales. The assumption made setting the cost of a shortage to the loss of unit profit may not be accurate for all items. Certain items have other products that can be easily substituted in their place, meaning that a lost sale for one item creates a sale for another item. Economists may further examine the effect of substitutions as they relate to superior and inferior goods, however for the purpose of the model, all it means is the cost of a unit shortage may not be as high as our assumption. This will be examined further in the sensitivity analysis.

The items used for the model also had the two highest holding costs of any of the eleven original items. Had the holding cost been lower for these items, the expectation is that the slope of the excess cost curve in Figure 4.2 would be lower, making it cheaper to hold more excess inventory. Holding costs are usually set by the cost of borrowing funds and are typically hard to change significantly without a major policy change from the Federal Reserve therefore this will not be looked at in the sensitivity analysis. However, the more interesting question is how the holding cost relates to the unit profit loss. The two items analyzed had profit to holding cost ratios of 2.25 and 3.16 for item 1 and item 3 respectively. These are two of the lowest three ratios of the original eleven items. The other items have ratios that range from 2.99 all the way to 61.85 with a weighted average (in terms of sales units) of around 9.23. This would indicate that the items analyzed had less emphasis on keeping excess inventory than the average item. This is something that will be analyzed in the sensitivity analysis.

The optimization calculations did not take into account the possibility of a set service level. As mentioned previously, given a set service level the optimal order quantity can be found using Equation 2.6 from Chapter 2. However, this equation still won't take into account the possibility that a higher service level may produce a lower total cost when the cost of a missed sale greatly exceeds that of holding excess inventory. The effect of service level will be analyzed in the sensitivity analysis as well.

One area that should be considered for all companies that rely on shipments from suppliers is the lead time. Companies can often work with their suppliers to either reduce the lead time or at least reduce the variability. Since this particular client owns its own distribution centers, the lead time for this model was assumed to be constant. However, there may be leeway with the length of the lead time through upstream optimization. This measure will be explored in the sensitivity analysis in order to determine if such optimization efforts are worthwhile.

4.2 Sensitivity Analysis

As outlined in the previous section, this section will look at how sensitive the inventory model is to changes in several of the inputs and underlying assumptions of the system. The first assumption to examine is whether or not the cost of a shortage is exactly equal to the profit lost from the missed sale. Using the model for both items at location A, Table 4.1 shows the effect that 25 and 50% reductions and increases in the shortage cost have on the order quantity and the total cost.

Table 4.1 Effect of Shortage Cost on Order Quantity and Total Cost for Location A.

Item	Factor	Cost of Shortage	Optimal Order Quantity	Total Cost	Percent Change from Baseline Cost
1	0.50	$ 13.30	10	$ 67.19	-11.3%
	0.75	$ 19.94	10	$ 71.49	-5.7%
	1.00	$ 26.59	10	$ 75.79	0.0%
	1.25	$ 33.24	11	$ 77.45	2.2%
	1.50	$ 39.89	11	$ 79.02	4.3%
3	0.50	$ 18.01	6	$ 41.54	-9.9%
	0.75	$ 27.01	7	$ 44.84	-2.7%
	1.00	$ 36.01	7	$ 46.08	0.0%
	1.25	$ 45.01	7	$ 47.32	2.7%
	1.50	$ 54.02	7	$ 48.56	5.4%

It appears as though the cost of shortage has a very limited effect on both the order quantity and the total cost. As expected, a decrease in the cost of shortage resulted in a lower total cost, and an increase led to an increase in total cost. In both cases a 50% reduction in this cost only led to approximately 10% savings. When the estimate for shortage cost increases by 50% the total costs increased by about 5%. The order quantities did not change very much for either item. The optimal order quantity for item 1 increased to 11 for both a 25 and 50% increase in the shortage cost. The order quantity for item 3 decreased to 6 when there was a 50% reduction in shortage cost. These changes likely had less to do with the fact that an overestimate or underestimate has a larger impact, and more to do with whether the original order quantity was on the cusp of changing to a lower or higher value. The argument that the shortage cost is overestimated when set equal to the profit lost from a sale because substitute goods exist may seem like a valid concern, especially when applying the total cost savings across thousands of SKUs. The argument that there is an additional cost that should be tacked on to the lost profit because of loss of customer loyalty could also be a legitimate concern when considering the total cost of managing inventory. When deciding how to establish to shortage cost, it would be wise to consider both of these scenarios and factor the lost profit figure accordingly.

The next measure to look at is the relationship between the yearly holding cost and the unit cost of shortage. As mentioned, the items selected had the lowest ratio of unit profit to holding cost of all the items in the data. To analyze the effect this has on the model, the unit profit level will be changed to create the desired profit to holding cost ratios for item 1 at location A. The results of this analysis is shown in Table 4.2.

Table 4.2 Effect of Profit to Holding Cost Ratio on Order Quantity and Total Cost

Item	Unit Profit	Holding Cost	Profit to Holding Cost Ratio	Optimal Order Quantity	Total Cost	Percent Change from Baseline Cost
1	$5.92	$11.84	0.50	9	$57.28	-24.4%
	$11.84	$11.84	1.00	10	$66.25	-12.6%
	$26.59	$11.84	2.25	10	$75.79	0.0%
	$47.36	$11.84	4.00	11	$80.79	6.6%
	$94.72	$11.84	8.00	12	$88.05	16.2%
	$109.28	$11.84	9.23	12	$89.81	18.5%
	$142.08	$11.84	12.00	12	$91.72	21.0%
	$732.30	$11.84	61.85	14	$107.41	41.7%

As with the shortage cost analysis, at first glance it looks as though the profit to holding cost ratio has a limited effect on the order quantity and the total cost. The selected ratios were chosen to represent hypothetical and real values of the ratio. The profit to holding cost ratio for item 1 is the lowest of all eleven items at 2.25, so clearly the ratios of 0.50 and 1.00 may not be possible, but they're worth calculating to examine the effect. In both cases the total cost decreased which is expected. In the extreme case for the 0.50 ratio there was a 24.4% decrease in total cost and a drop from 10 to 9 for the optimal order quantity. This 24.4% drop seems significant but given that there was nearly an 80% reduction in unit profit lost, the drop is not that severe. On the opposite extreme, when the ratio was equivalent to the maximum observed value from the data of 61.85, the total cost rose 41.7% and the order quantity went from 10 to 14. This 41.7% also pales in comparison to the 2654% increase in unit profit lost. The clear takeaway from this analysis is that as the ratio of profit to holding cost increases, the optimal order quantity, and therefore the service level and safety stock, all increase to protect against a missed sale. To check the validity of these calculations, the unit profit was also held constant at baseline while the holding cost varied to reach the desired ratios. These results looked nearly identical to the results in Table 4.2, with the exception of some slight variation as the holding cost approached zero.

The next measure that will be analyzed is the effect of the service level. The higher the service level the more priority is placed on keeping excess stock. A set service level higher than the optimal value may be put in place through a corporate strategy with the intention of raising customer service. Likewise a lower than optimal service level may be set to lower the cost of excess inventory. In the calculations from Chapter 3, situations were considered for when the service level was set, and when the service level was optimized as a result of the cost function. Using the Equation 2.6 from Chapter 2 the service level for item 1 from store A is examined in Table 4.3.

Table 4.3 Effect of a Set Service Level on the Order Quantity and Total Cost.

Set Service Level	Optimal Order Quantity	Total Cost	Actual Service Level
80.0%	7	$218.88	0.81879
85.0%	8	$131.90	0.89766
90.0%	9	$90.21	0.94801
94.0%	9	$90.21	0.94801
95.0%	10	$75.79	0.97634
97.5%	10	$75.79	0.97634
98.0%	11	$75.87	0.99038
99.0%	11	$75.87	0.99038
99.5%	12	$82.76	0.99652
99.9%	14	$103.18	0.99968

From this table it can be seen that setting a service level can be misleading for certain items. For item 1 at store A when the set service level is 95% or 97.5% the total cost is minimized. However, if another service level is chosen due to corporate strategy or another reason, the cost is not minimized, and will cost the company money in the long run. Looking at the order quantities for a service level of 94% versus 95% will illustrate how this may happen. The difference between these two values is very minimal and 1% could be the difference between what is the chosen service level, and the optimal service level. For these values there is a difference in the optimal order quantity and a significant difference in the total cost of the policy. In a similar fashion, corporate strategy could set the service level higher than the optimum, which would cause the company to spend more money than necessary to keep excess inventory. The takeaway from this analysis would be to set the corporate strategy for items with the optimum service level in mind. If there is a significant reason to set the service level higher or lower than the optimum it would have to outweigh the additional cost of modifying the service level.

The last variable in the model that will be analyzed is the lead time. The lead time in the base stock system is the only variable that increases the uncertainty of the demand. From Equation 2.6 it is apparent the lead time has an effect on both the mean and the standard deviation of the normal distribution used to describe the demand. If the lead time has enough of an effect on the total cost of the policy then it may be worth trying to work with the upstream suppliers to reduce the lead time. The effect of the reducing and increasing the lead time for item 1 at store A by 25 and 50% is shown in Table 4.4.

Table 4.4 Effect of Lead Time on Order Quantity and Total Cost.

Factor	Lead Time	DLTR Mean	DLTR Standard Deviation	Optimal Order Quantity	Total Cost	Percent of Baseline Cost
0.50	0.50	3.3396	2.4221	9	$64.40	-15.03%
0.75	0.75	3.8962	2.6162	10	$70.24	-7.32%
1.00	1.00	4.4528	2.7968	10	$75.79	0.00%
1.25	1.25	5.0094	2.9664	11	$79.33	4.67%
1.50	1.50	5.5660	3.1269	12	$83.14	9.70%

The lead time seems to have a fairly significant effect on the optimal order quantity. This is probably because the lead time directly effects both the mean and standard deviation of the normal distribution used to model the demand. When the lead time is half the given value, the total cost is reduced by 15.03%. Similarly, when the lead time is increased by 50% the total cost goes up by 9.7%. This is expected because as the lead time decreases the lead time demand becomes smaller along with the associated variation in the demand during that time. When the lead time is smaller, less excess product is required to protect against a stock out, causing the optimal order quantity to be lower. Since the order quantity is lower, the average excess inventory is also lower. Both of these factors combine to lower the total cost of managing the inventory. The takeaway from this analysis indicates that if the lead time can be reduced significantly enough to justify the upstream costs of reducing that lead time, then it is worthwhile.

Chapter 5

Conclusions and Recommendations

5.1 Conclusions

The objective of this thesis was to develop a model to manage inventory at a large national retailer. By researching the available inventory policies used in industry, a framework for solving this problem was laid out. The data from the retailer was closely examined, first to develop an ABC classification scheme for the items which prioritized the largest item holding costs. The weekly demand for two items was examined at two different locations over the course of one year. This demand was fit to a normal distribution that described the average weekly demand and standard deviation for each item. Using information from the client about the way the supply chain for the retail stores is managed, a base stock policy was selected to simplify the management of inventory at the store level. With the policy and demand distribution in place, the optimal order-up-to levels were calculated using a cost minimizing equation in excel solver. The results of this optimization calculated the optimal order level for item 1 at location A to be 10 units with a service level of 97.634% and a total annual cost of $75.79. The order level for item 1 at location B was also determined to be 10 units, with a service level of 98.178% and a total annual cost of $73.45. The optimum order level for item 3 at location A is 7 units with a service level of 99.238% and a total annual cost of $46.08. Lastly, the order level for item 3 at location B was determined to be 12 units with a service level of 98.734% and a total annual cost of $84.27.

Through a sensitivity analysis certain assumptions and inputs of the model were analyzed for their effects on the output. The cost of shortage, unit profit to holding cost ratio, service level, and lead time were all analyzed at values lower and higher than the original values in the data. The cost of shortage was originally estimated to be equal to the profit lost from a missed sale. The sensitivity analysis revealed that altering this value won't have a significant effect on the order quantity, but can have a significant effect on the total annual cost estimate. The effect on the total cost is especially true if the extra savings or expenditures is applied across thousands of SKUs. The unit profit to holding cost ratio sensitivity analysis revealed that as the ratio of profit to holding cost ratio increases, the optimal order quantity, service level and safety stock all increase at a moderate rate. This ratio should be considered when applying this model to other items, however since the ratio is difficult to change for a given item there should not be much emphasis placed on it. The service level sensitivity analysis showed the potential cost implications of setting a service level significantly higher and lower than the optimal level. As the service level increased, so did the order quantity, thus spending more to keep excess inventory on hand. When there is a significant reason to alter the service level, either above or below the optimal level, it should be considered alongside the additional cost. Finally the lead time sensitivity analysis showed that decreasing the lead time would likely lower the optimal order level causing the total annual cost to decrease. The opposite effect occurs when the lead time increases. When deciding whether or not to try and lower the lead time, the cost of doing so should be compared with the savings potential for the particular reduction.

5.2 Recommendations

The conclusions presented in the previous section are general conclusions that can be applied to any company using a similar inventory management policy. Specific recommendations for the client can be made from these general conclusions. When applying this model to the other items across the retail chain, the client should first consider the specific item demand. This model was developed for items that showed consistent demand throughout the year and modeled that demand as normally distributed. Before applying this model to other items, the demand for those items should be evaluated to see if the normal assumption is valid. If the assumption holds then the model can be applied.

When setting up the optimization, the first parameter to consider is the unit shortage cost for each item. As identified in Section 4.2, this value can significantly affect the cost estimations when it is applied to thousands of SKUs as it would be for the client. If the item in question has several substitute goods available, the shortage cost should be estimated to be lower than the unit profit lost for a missed sale. If the item has no substitutes, the unit profit is probably a good estimate. If customer loyalty is in jeopardy when the item is out of stock, the shortage cost should be estimated higher than the unit profit to put more emphasis on the cost of a missed sale.

The next consideration is whether or not to increase or decrease the optimum service level. For most items, the optimum service level should be used because it is set to provide the appropriate amount of protection against a stock out. As stated, there must be a reason to do this that outweighs the additional annual cost. If the product is a store brand product that is trying to build a consistent customer base, then perhaps it is worth increasing the service level. If a product is being replaced by a new version, then it may make sense to lower the service level for the older version to prepare for the switch.

The final measure to consider for the clients stores is the product lead time. As mentioned, this value may have an implication on the optimum order level. From the original data in Table 3.1, the lead time values are different for different products. In the model the largest lead time was selected for the calculations. This was not a problem for the items analyzed because the lead time for the two items were similar enough where this wasn’t a major factor. But when planning for an order with hundreds or thousands of different SKUs, the longest lead time may be significantly larger than the shortest. The effect of increasing the lead time for several items will cause the order quantities and the annual costs to increase. If the high lead times can be reduced through an optimization effort at the distribution centers, average inventory levels for all items at the stores can be reduced. This would likely save on annual inventory management costs across the chain, depending on how costly the upstream optimization effort is.

Appendix A

Raw Sales Data

Table A.1 Total Sales Quantities During Collection Period.

Item	Total Sales Quantity		
	Location A	Location B	Total
1	117	88	205
2	835	3292	4127
3	46	94	140
4	2	11	13
5	12	31	43
6	12	282	294
7	0	5	5
8	0	13	13
9	46	243	289
10	19	23	42
11	9	15	24

Table A.2 Weekly Sales Quantities for Items 1 & 3.

Week	Item 1		Item 3		Week	Item 1		Item 3	
	Store A	Store B	Store A	Store B		Store A	Store B	Store A	Store B
7/7/2012	3	2	1	2	1/12/2013	0	2	0	5
7/14/2012	1	-1	0	1	1/19/2013	0	0	0	3
7/21/2012	0	0	1	1	1/26/2013	1	0	0	2
7/28/2012	3	2	2	0	2/2/2013	5	2	1	4
8/4/2012	5	2	2	0	2/9/2013	1	4	0	1
8/11/2012	1	1	1	1	2/16/2013	0	5	0	2
8/18/2012	1	1	3	0	2/23/2013	1	0	1	0
8/25/2012	1	0	1	1	3/2/2013	3	2	1	1
9/1/2012	1	5	0	1	3/9/2013	8	0	3	1
9/8/2012	2	2	1	0	3/16/2013	2	6	0	1
9/15/2012	2	3	2	3	3/23/2013	0	2	2	3
9/22/2012	0	1	-1	0	3/30/2013	1	1	3	5
9/29/2012	0	1	1	3	4/6/2013	2	2	1	3
10/6/2012	7	1	0	2	4/13/2013	1	3	0	0
10/13/2012	1	1	1	1	4/20/2013	0	3	0	2
10/20/2012	1	0	4	0	4/27/2013	4	0	2	1
10/27/2012	2	1	0	2	5/4/2013	4	4	0	1
11/3/2012	5	4	0	1	5/11/2013	1	0	3	2
11/10/2012	2	0	1	0	5/18/2013	4	0	1	2
11/17/2012	3	0	-1	0	5/25/2013	1	2	0	3
11/24/2012	2	1	1	0	6/1/2013	5	1	0	5
12/1/2012	5	1	1	0	6/8/2013	4	3	0	4
12/8/2012	1	1	0	1	6/15/2013	0	3	3	10
12/15/2012	5	3	-1	2	6/22/2013	3	5	2	0
12/22/2012	-1	0	2	4	6/29/2013	5	1	0	5
12/29/2012	5	4	1	0	7/6/2013	2	0	1	2
1/5/2013	1	1	-1	0					

Appendix B

Statistical Evaluation

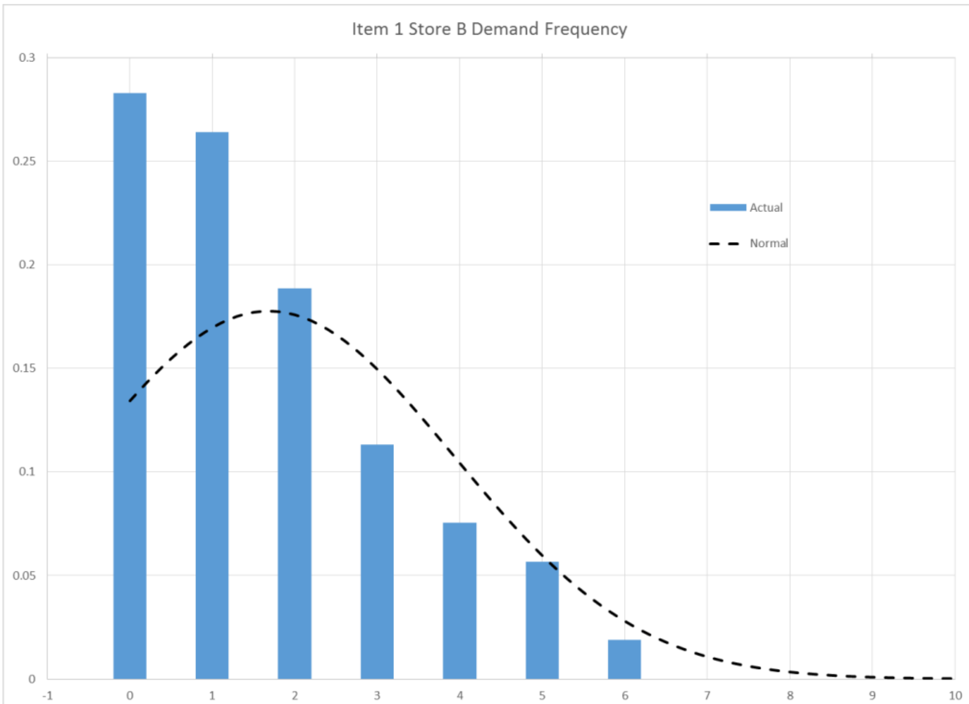

Figure B.1 Histogram of Item 1 Store B Demand Frequency

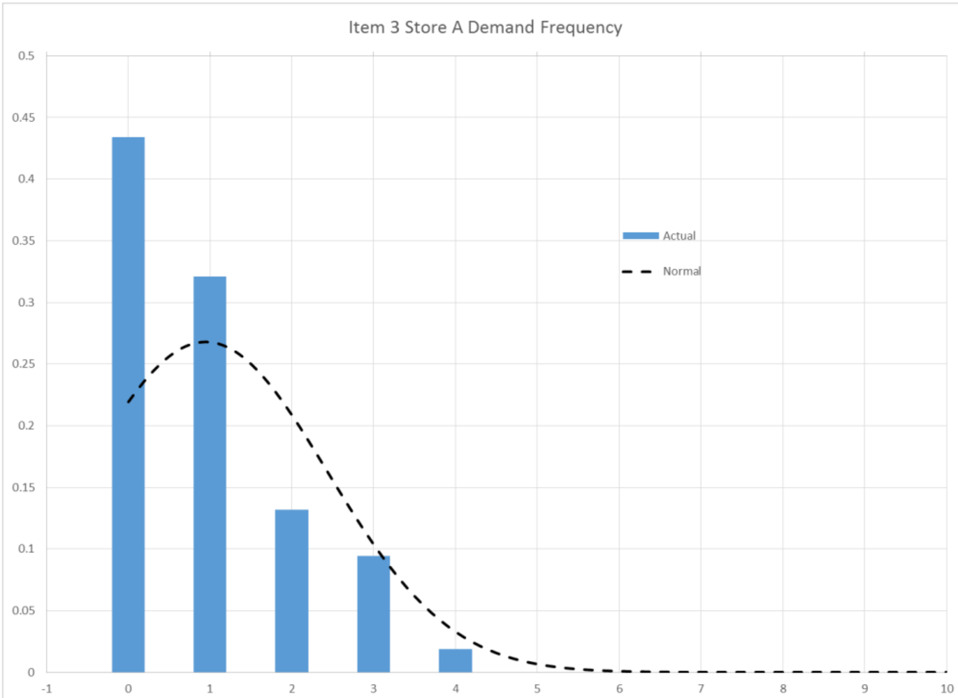

Figure B.2 Histogram of Item 3 Store A Demand Frequency

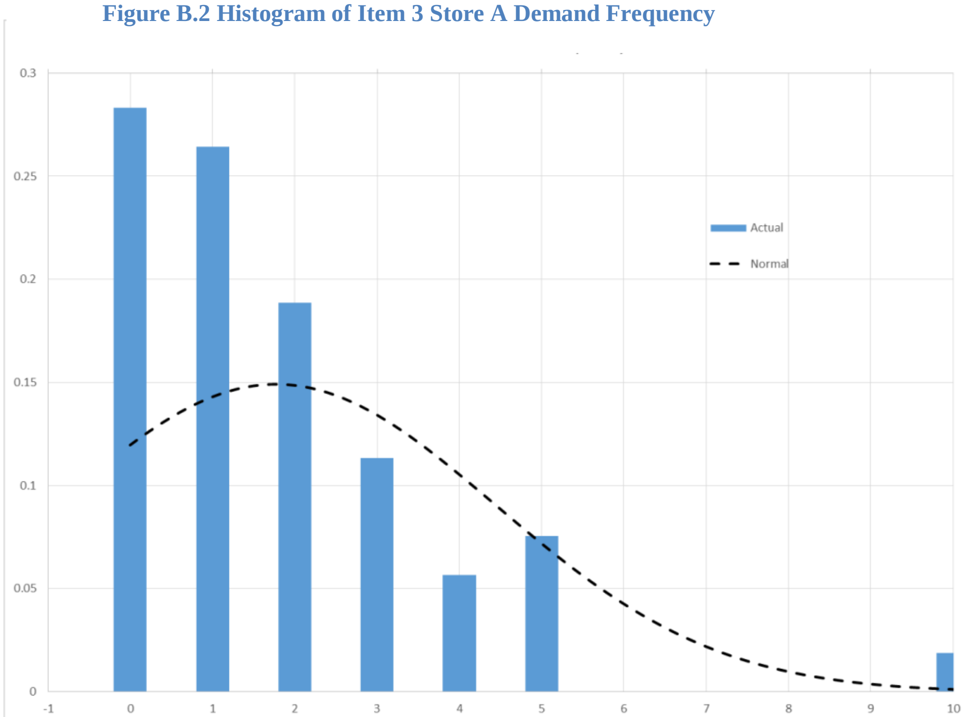

Figure B.3 Histogram of Item 3 Store B Demand Frequency

BIBLIOGRAPHY

Berling, P. (2008). Holding Cost Determination: An Activity-based Cost Approach. *International Journal of Production Economics*, *112*(2), 829-840. Retrieved from http://www.sciencedirect.com/science/article/pii/S0925527307002605

Feature report: Retail & inventory management systems. (1994, Nov 16). *Computer Dealer News, 10,* 31-48. Retrieved from http://search.proquest.com/docview/202749320?accountid=13158

Femia, J. V., & Marshall, A. J. (2012). Introduction. *Vilfredo Pareto beyond disciplinary boundaries* (Online-Ausg. ed., pp. 3-4). Farnham: Ashgate.

Haggerty, R. (1999, November 19). Numerical Integration Using Excel. *GEO 487: Introduction to Hydrology*. Retrieved November 1, 2013, from http://people.oregonstate.edu/~haggertr/487/integrate.htm

Jensen, P. A., & Bard, J. F. (2003). Inventory Theory [Supplemental Material]. *Operations research: models and methods* (pp. 1-55). Hoboken, N.J.: Wiley. Retrieved from http://www.me.utexas.edu/~jensen%20/ORMM/supplements/units/inventory/inventory.pdf

Mondaq Business Briefing. (January 25, 2010). United Kingdom: Emerging From The Downturn - Global Powers Of Retailing 2010., Retrieved from www.lexisnexis.com/hottopics/lnacademic

Plunkett, Jack W., Plunkett, M. B., Steinberg, J. S., Faulk, J., & Snider, I. J. (2012). Introduction to the Retail Industry.*Retail Industry*. Retrieved November 5, 2013, from http://www.plunkettresearchonline.com

Ravindran, A. R., & Warsing, D. P. (2013). Inventory Management Methods and Models. *Supply chain engineering: models and applications* (pp. 95-171). Boca Raton, FL: Taylor & Francis.

Teunter, R. H., Babai, M. Z. and Syntetos, A. A. (2010), ABC Classification: Service Levels and Inventory Costs. *Production and Operations Management, 19*: 343–352. Retrieved from http://search.proquest.com/docview/356625322?accountid=13158

ACADEMIC VITA

Patrick Thomas Lucas
18 Hawthorne St.
Woburn MA 01801
plucas005@gmail.com

Education

The Pennsylvania State University
The Schreyer Honors College
Bachelor of Science in Industrial Engineering (Fall 2013)

Thesis Title: Use of Inventory Management for Large General Item Retailers
Thesis Supervisor: Dr. A. Ravi Ravindran

Honors and Awards

- Julia and Brent Beabout Honors Scholarship (Fall 2013)
- IE Outstanding Senior Award Finalist (Spring 2013)
- Life Technologies Honors Scholarship in Industrial Engineering at Penn State (Fall 2012)
- Harold and Inge Marcus Industrial and Manufacturing Engineering Undergraduate Scholarship (Fall 2011)
- Dean's List (2009-2013)

Association Memberships/Activities

- Alpha Pi Mu Industrial Engineering Honor Society

Professional Experience

Capital One Financial
Operations Analyst Intern – Network Strategy
Summer 2013

Accenture
Analyst Intern – Operations Workforce Optimization
Fall 2011 & Summer 2012