

DETERMINING OPTIMAL INVENTORY LEVELS FOR ITEMS NEARING THE END OF A PRODUCTION RUN

A Thesis presented to
The Faculty of the Graduate School
At the University of Missouri-Columbia

In Partial Fulfillment
Of the Requirements for the Degree
Master of Science, Industrial Engineering

by
LUCAS CATO
Dr. Ronald McGarvey, Thesis Supervisor
DECEMBER 2020

The undersigned, appointed by the dean of the Graduate School, have examined the thesis entitled:

**DETERMINING OPTIMAL INVENTORY LEVELS FOR ITEMS
NEARING THE END OF A PRODUCTION RUN**

presented by **Lucas Cato**,

a candidate for the degree of Master of Science and hereby confirm that, in their opinion,
is worthy of acceptance.

Dr. Ronald McGarvey

Dr. James Noble

Dr. Anthony Ross

ACKNOWLEDGEMENTS

First, I would like to thank Dr. Ronald McGarvey and Dr. James Noble and for their continued support and guidance throughout the completion of my degree program. They have continuously demonstrated a willingness to share their expertise with me, inspiring my curiosity and motivating me to do more. It has been a great honor to work with them throughout the completion of this research and I truly believe that I could not have accomplished this endeavor without them. I would like to give a sincere thanks to Dr. Anthony Ross for his flexibility and understanding throughout the process of defending this thesis. I would also like to express my sincere gratitude to all faculty and staff of the Industrial Engineering department here at the University of Missouri-Columbia and especially Dr. Cerry Klein for motivating me to complete a graduate level degree.

Lastly, I want to thank the Center for Excellence in Logistics and Distribution (CELDi) research team for their support and encouragement. The CELDi team introduced me to supply chain related research, leaving me with a passion for solving the complex and integrated problems within this field. I would like to give a special shout out to Randolph Bradley for acting as my industry mentor over the past three semesters. Randolph has continuously made himself available for help no matter the time or day. His constructive criticism and feedback has had an immeasurable impact on my professional development.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	ii
LIST OF TABLES	v
LIST OF FIGURES	vi
ABSTRACT	vii
Chapter 1: Introduction	1
1.1. Aerospace Supply Chain Structure	2
1.2. An Example of the Aerospace Supply Chain Structure	9
Chapter 2: Literature Review	14
2.1. Newsvendor Model Research	14
2.2. (s, S) Inventory Policies Research	14
2.3. Distribution Fitting	15
2.3.1. Gamma Distribution	16
2.4. Shortages & Backorders	20
Chapter 3: Methodology	21
3.1. Parameter Definitions	21
3.1.1. Unit Ordering Cost (c)	21
3.1.2. Storage Cost	21
3.1.3. Salvage Value	21
3.1.4. Holding Cost (h)	22
3.1.5. Penalty Cost (p)	22
3.1.6. Set Up Cost (k)	23
3.1.7. Initial Inventory (x)	23
3.2. Lead Time Scaling	23
3.3. News Vendor Model Formulation	24
3.4. (s, S) Model Formulations	27
3.4.1. (s, S) Equation Based	27
3.4.1.1. (s, S) Equation Based with Normal Distribution	31
3.4.1.2. (s, S) Equation Based with Gamma Distribution	32
3.4.2. (s, S) Algorithm Based	34
3.4.2.1. (s, S) Algorithm Based Implementation Challenges	36
3.5. Computational Methods	45

Chapter 4: Results & Analysis	47
4.1. Data	47
4.2. Example Problems	47
4.2.1. Newsvendor Example	50
4.2.2. A Note on (s, S) Policies	52
4.2.3. (s, S) Equation Example	54
4.2.4. (s, S) Algorithm Example	56
4.3. Cost Benefit Analysis	57
Chapter 5: Conclusion & Future Work	67
5.1. Conclusion	67
5.2. Future Work	68
BIBLIOGRAPHY	69

List of Tables

Table	Page
1. Safety Stock & Lead Time Guidelines	8
2. Scrap, Loss & Rework Example	12
3. (s, S) Algorithm Based Example Problems	43
4. Example Dataset	47
5. (s, S) Example Problem Parameters	48
6. (s, S) Example Problem Cost Calculations	49
7. (s, S) Newsvendor Example Problem	52
8. Gamma Function Example	53
9. (s, S) Equation Example	55
10. (s, S) Algorithm Example	57
11. Cost Benefit Analysis Overview	60
12. Proportion of Total Cost Accounted for by Additional Costs	62

List of Figures

Figure	Page
1. Traditional Supply Chain Structure for Aircraft Manufacturing	3
2. Redesigned Supply Chain Structure for Dreamliner Platform	3
3. The Bullwhip Effect	4
4. Push vs. Pull vs. CONWIP	6
5. MRP system	7
6. Example Aerospace Supply Chain Structure	9
7. Gamma Density Function Beta = 1	17
8. Gamma Density Function Beta = 2	17
9. Comparison of Poisson and Gamma Distributions	19
10. Comparison of Binomial and Gamma Distributions	19
11. Comparison of Shifted Lognormal and Shifted Gamma Distributions	19
12. Graph of One Period Expected Holding Cost Vs. Inventory Position	37
13. Graph of One Period Expected Shortage Cost Vs. Inventory Position	38
14. Graph of $L(y)$ Vs. Inventory Position	38
15. Graph of $G(y)$ Vs. Inventory Position	40
16. Graph of $G(S)$ and c^0 Vs. Inventory Position	41
17. MATLAB Script Example	46
18. Additional Cost Proportions-Normal Distribution	62
19. Additional Cost Proportions-Gamma Distribution	63
20. T-test on Newsvendor and Baseline for Normal Distribution	64
21. T-test on Newsvendor and Baseline for Gamma Distribution	65

22. T-test on Gamma and Normal Newsvendor Models 66

ABSTRACT

Optimal inventory policies determined for items in steady state production are no longer optimal when reaching the end of production. This is due to the obvious fact that during this time, production is no longer in steady state. At the end of a production run, steady state inventory policies can lead to excess costs, as on hand and due-in inventory is no longer need as there is no following period's demand. In this thesis, a newsvendor inventory optimization model which considers salvage value and initial inventory level along with two alternative (s, S) model formulations are tailored to fit items nearing the end of a production run. One of the (s, S) inventory models is modified to include a salvage value and reduce computation time. The three models are demonstrated on a twenty-item example problem and the newsvendor model is selected as the best application for items one lead-time period from the end of production. The cost related benefits of the alternative inventory policies generated by the newsvendor model are analyzed using a simulation-based approach. Although the ideas and analysis presented here were developed and tailored to fit the aerospace industry, the mathematical models can be extended to fit a variety of different applications.

Chapter 1

Introduction

The objective of this thesis is to identify improved methods for determining inventory levels for items nearing the end of a production run. Previously calculated steady state inventory levels will no longer be optimal solutions during this time frame. This is because the assumption of steady state production is not valid as production shifts toward a stop. Thus, maintaining steady state inventory levels during this time frame will result in excess costs to the supply chain. As a product reaches the end of a production run, it is critical to determine inventory levels that accurately fit demand, as excess raw material or parts will directly result in inventory at the end of the period. At the same time, having too few parts at the end of a production run will still result in a shortage (as it would at any other point in time), causing the part to be backordered or worse, delaying delivery of a complex end item such as an aircraft. In addition, suppliers may have little incentive to restart their now idling production lines given that the quantity of shortfall is a relatively small number of parts.

A company may determine inventory levels during this critical time by drawing down safety stock levels, resulting in very little or in often cases no room for error. In the aerospace industry, a stockout, or shortage, will result in the delayed delivery of a high value aircraft, causing excessive holding cost and a potential penalty cost for a missed delivery date. This leaves room for significant inventory cost savings throughout these time periods by balancing the cost of excess against the penalty cost of shortages. Given the current COVID-19 impacts on the aerospace industry, it is of upmost importance to minimize costs and stay as lean as possible. Balancing these costs provides the motivation for this thesis topic.

Inventory optimization has already proven to result in substantial cost savings in many different industries including the aerospace industry in which this study will take place. This thesis will take an in-depth look into this unique inventory optimization sub-problem, exploring several ideas and mathematical models in search of an optimal or near optimal solution. A cost benefit analysis will be conducted to estimate the monetary value of such improvements to these ends of production inventory level policies.

1.1 Aerospace Supply Chain Structure

The aerospace industry makes for an interesting application of this problem. The first thing to note is the sheer complexity of aerospace supply chains. To understand the complexities involved in aerospace related supply chains, we will look at a case study contrasting the Boeing 787 Dreamliner and 737 supply chain structures. To significantly reduce development cost and time for the 787 Dreamliner an unconventional supply chain design was used instead of the traditional supply chain structure utilized for the 737 platform (Tang & Zimmerman 2009). The Dreamliner supply chain structure utilized a tiered system in which Boeing would partner with roughly 50 tier-1 suppliers who would assemble parts/sub-assemblies produced by tier-2 suppliers, as opposed to the traditional structure in which Boeing would assemble parts/sub-assembly from thousands of different suppliers (Tang & Zimmerman 2009). The traditional supply chain structure can be seen in Figure 1 while the redesigned structure is seen in Figure 2.

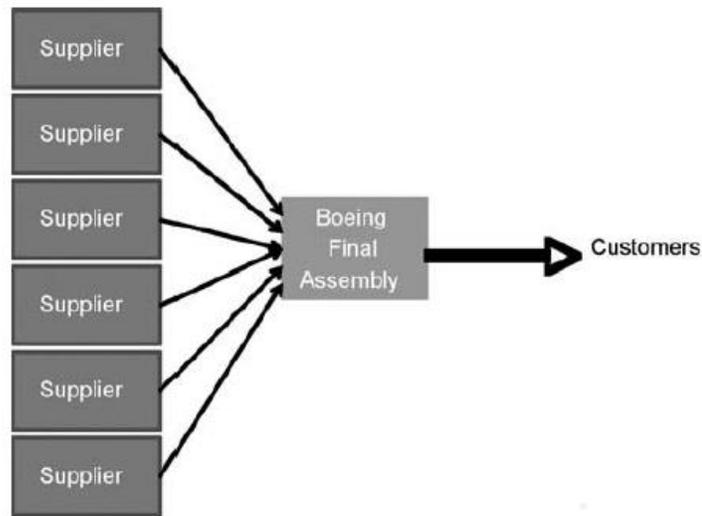


Figure 1. Traditional Supply Chain Structure for Aircraft Manufacturing

Source: Christopher S. Tang and Joshua D. Zimmerman, A traditional supply chain for airplane manufacturing

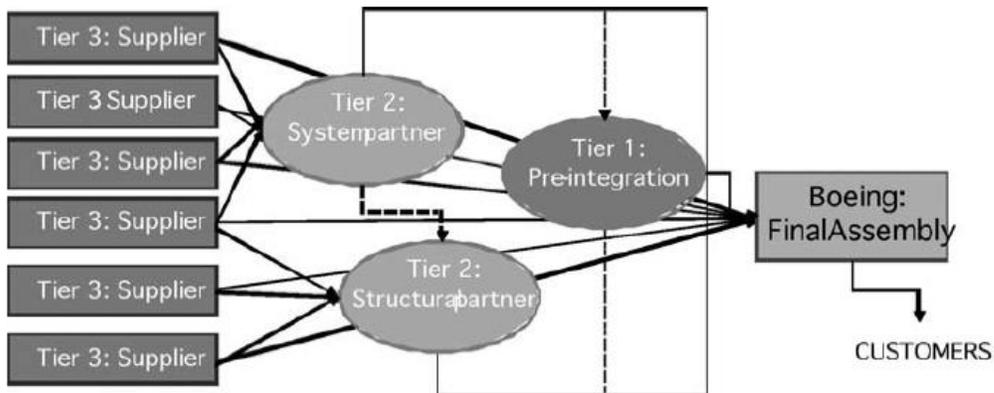


Figure 2. Redesigned Supply Chain Structure for Dreamliner Platform

Source: Christopher S. Tang and Joshua D. Zimmerman, Redesigned supply chain for the Dreamliner program

The redesigned structure increases outsourcing percentage to 70% for the Dreamliner as opposed to 35-50% in the traditional supply chain design used by the 737 platform (Tang & Zimmerman 2009). Although the redesigned structure reduces the complexities directly related to Boeing by

increasing the outsourcing percentage and sharing risk sharing between stages of suppliers, this redesigned structure increases the complexity of the overall supply chain (Tang & Zimmerman 2009). This increased complexity for the overall supply chain causes the supplier risk for Boeing to increase dramatically. As a relatively small interruption in a tier-3 supplier can now be compounded by tier-2 and tier-1 suppliers, causing a bullwhip effect that leaves the final assembly of the aircraft severely disrupted by a seemingly small disruption from a high tier supplier. For readers unfamiliar with the bullwhip effect, this is a phenomenon referred to as the variability in demand order quantities in the supply chain being amplified as they move up the supply chain (Lee, Padmanabhan, & Seungjin 1997).

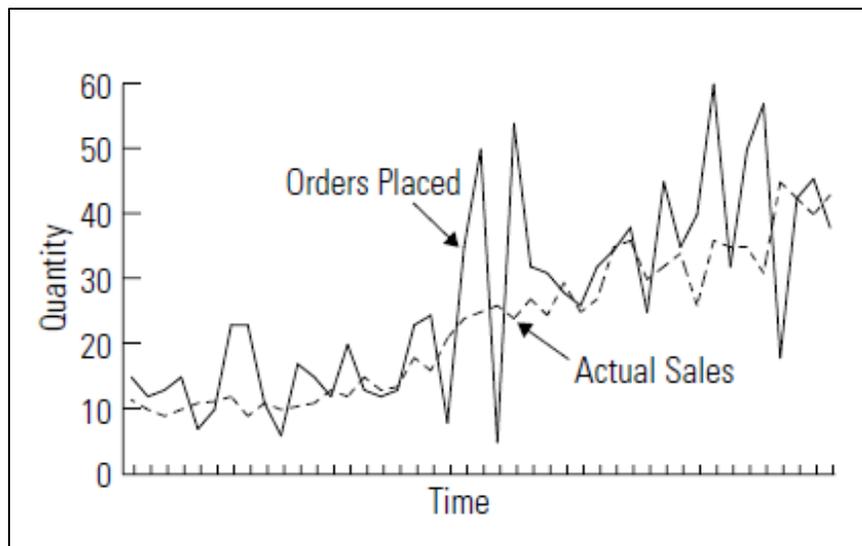


Figure 3. The Bullwhip Effect

Source: Hau L. Lee, V. Padmanabhan, and Seungjin Whang, Higher Variability in Orders from Dealer to Manufacturer than Actual Sales

However, a similar phenomenon can happen in the opposite direction, from suppliers to consumers. Another thing to note from this case study is the overall number of suppliers in which inventory levels must be managed to build a single aircraft platform. Keep in mind that often

multiple parts are procured from a single supplier, implying that thousands of inventories polices must be accurately determined at the time that a single product or aircraft reaches the end of a production run.

In the aerospace industry the products manufactured and sold are of very high monetary value and aerospace firms either don't have the available assets or it doesn't make sense financially to manufacture them in advance of another party agreeing to purchase them. Therefore, it is common in the aerospace industry for products to be sold based on contracts between buyers and sellers, with a predetermined number agreed upon. This implies that there is little demand uncertainty in the aerospace industry. Extending this logic, this means that most of the variation in demand for sub-assemblies and parts comes primarily from scrap, loss, and rework in the manufacturing processes, as opposed to inaccurate demand forecasts, as is the case in many other industries. This will be a key idea throughout the thesis. Another caveat to aerospace companies operating primarily off purchasing contracts is that when a part shortage occurs, the demand is backordered and fulfilled later. In other industries, it can be common for a shortage to result in a loss sale. However, this will not be the case in this analysis where all part shortages are assumed to result in the demand being backordered and fulfilled later. This will be discussed further later.

It is common to classify a production system as either a push, pull or a hybrid of the two. A push system is one in which a job is started on a start date that's based off of the established lead time while a pull system refers to a downstream work center pulling stock from previous operations as they need it (Spearman & Zazanis, 1992). The pull system makes it such that all work is being performed only to replenish the outgoing stock (Spearman & Zazanis, 1992). The pull system is

designed to cap the amount of work in progress in each system while the push system has no cap on the amount of inventory in the system (Hopp & Spearman 2004). A CONWIP system would be an example of a hybrid of both push and pull systems. In a CONWIP system, inventory is only pulled the front of the line and inventory is pushed in between the next pull (Spearman & Zazanis, 1992). The three different production systems can be seen below in Figure 4.

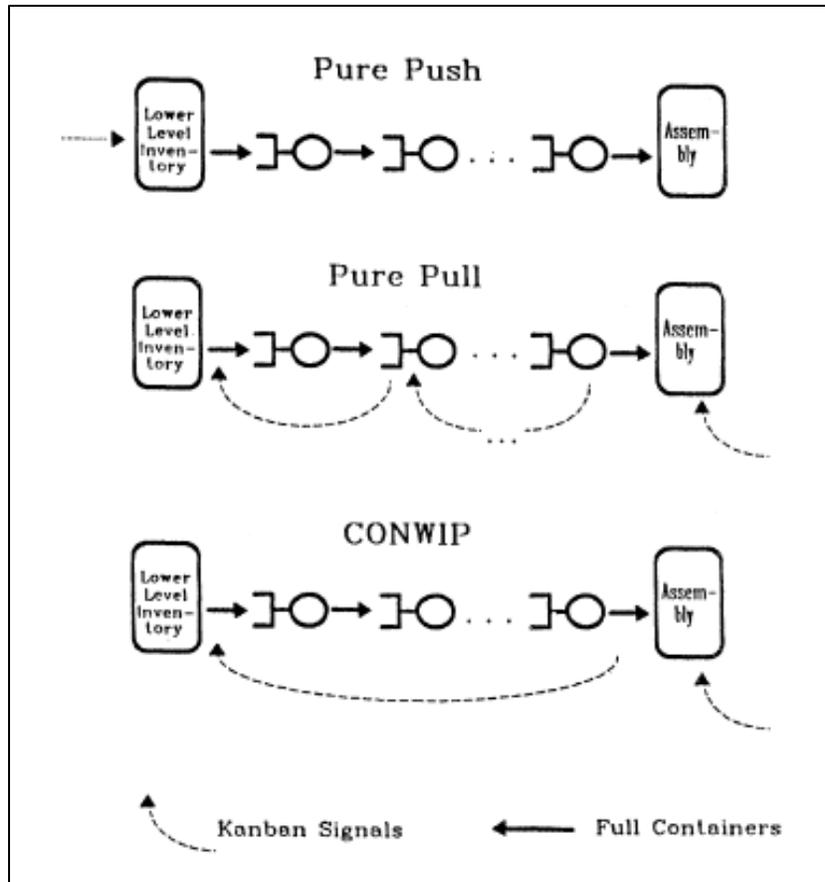


Figure 4. Push vs. Pull vs. CONWIP

Source: Mark L. Spearman and Michael A. Zazanis, Pure push, pure pull and constant WIP systems.

Due to the small demand uncertainty and high manufacturing complexity, aerospace products naturally fall into the push category of manufacturing systems. In fact, many aerospace companies use a Material Resource Planning (MRP) based push system.

The MRP systems considers a products bill of materials along with the corresponding lead times to create a plan for when jobs should start, or inventory gets released. The inputs and outputs of an MRP system can be seen below in Figure 5.

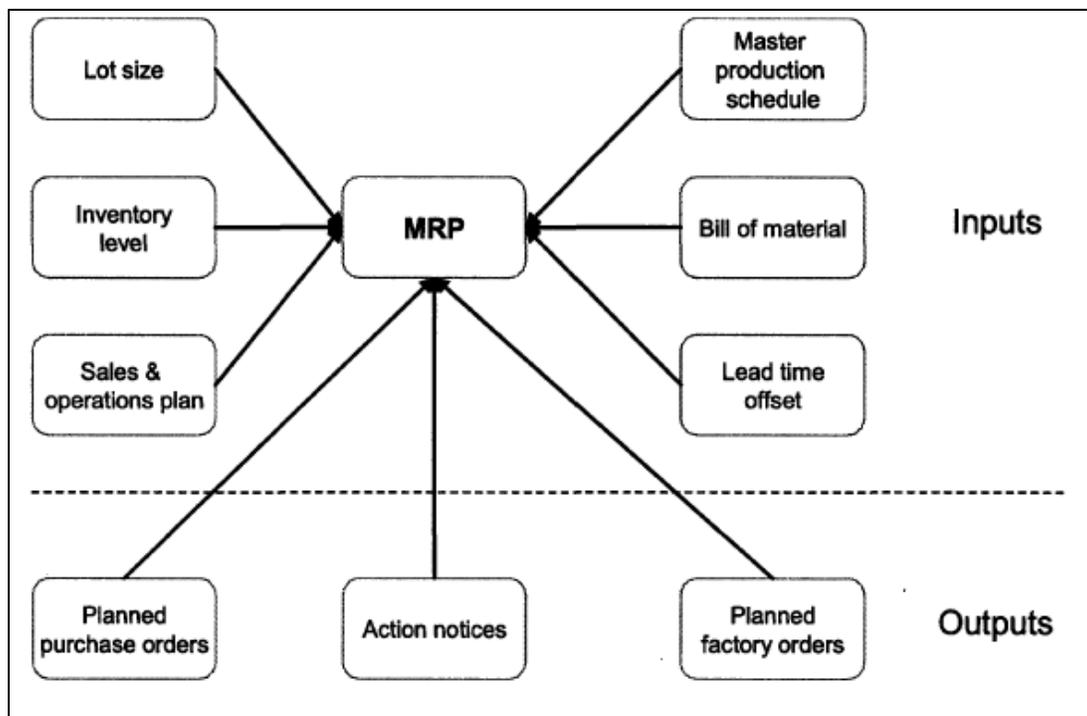


Figure 5. MRP system

Source: Joseph L. Aiello, Key Inputs and Outputs of an MRP System.

Because this is a push system, there is no cap on inventory or work in process. MRP systems assume that part lead-times are fixed and do not account for inventory already in the plant (Hopp & Spearman, 2001). However, the time for parts to travel through a manufacturing plant does in fact depend on the amount of inventory already in the plant, thus, the fixed lead time assumption is clearly not valid and sure to cause uncertainties (Hopp & Spearman, 2001). There are traditionally two ways to deal with uncertainty, by using safety stock inventory or by using safety

lead time. Safety stock refers to excess inventory strategically set to act as a buffer against this uncertainty, while safety lead time represents a time buffer as opposed to inventory, both are set with the goal to protect against variation or uncertainty (Silver, Pyke, & Peterson, 1998). Whybark and Williams (1976) analyze these safety buffers under the following four different instances (Mauro 2008):

1. Supply Timing: orders not received when scheduled
2. Supply quantity: orders received for more or less than planned
3. Demand timing: requirements shift from one period to another
4. Demand quantity: requirements for more or less than planned

The results from their analysis suggest that uncertainties involving timing are best buffered by safety lead-time, while those uncertainties involving quantity are best buffered against with safety stock (Whybark & Williams, 1976). Silver, Pyke, and Peterson (1998) also discuss this in their analysis and make the following recommendations seen in Table 1 below.

Safety stock	Safety lead-time
<ul style="list-style-type: none"> • In items with direct external usage • In items produced by a process with a significantly variable yield • In items produced at a bottleneck operation • In certain sub-assemblies used for a myriad of end-items 	<ul style="list-style-type: none"> • In raw materials

Table 1. Safety Stock & Lead Time Guidelines

Source: Joseph Mauro, Safety Stock and Lead Time Guidelines

Noting that as discussed earlier, we are interested in demand quantity variation with respect to scrap, loss, and rework, which also corresponds to items with a significantly variable yield,

confirming the suspicion that safety stock levels should be explored in the problem presented throughout this thesis.

1.3 An Example of the Aerospace Supply Chain Structure

The mathematical modeling and analysis portion of this thesis will be tailored to fit the aerospace supply chain structure. However, we expect that the models and ideas developed for the aerospace industry can also easily be adjusted for implementation across a variety of different applications. Although, it is important to note that the structure of supply chains outside of the aerospace industry may be significantly different. So, the need for adjustments to be made before switching over between applications is an important concept to be considered. For the remainder of this section we will investigate the typical structure of an aerospace supply chain and how it relates to our analysis.

Figure 6 below represents the typical supply chain structure of an example aerospace company.



Figure 6. Example Aerospace Supply Chain Structure

Looking at the structure we see that there are multiple suppliers from which an aerospace company procures raw materials and parts/sub-assemblies that feed production. Once the procured items reach the company they are then put into an inventory stock room for storage. The items will be stored until needed by a mechanic working on the shop floor, at which point they will be issued

from the stock room to the shop floor. The shop floor mechanic will attempt to install the item onto the aircraft, and one of four following things will happen:

1. The shop floor mechanic successfully installs the item on the aircraft
2. The shop floor mechanic loses the item while attempting to install it onto the aircraft
3. The shop floor mechanic will unsuccessfully install the item, resulting in the item being scrapped
4. The shop floor mechanic will unsuccessfully install the item, resulting in the item needing to be reworked

In the case that the first event happens and the shop floor mechanic successfully installs the item then the aircraft will move onto the next stage in production, eventually being delivered to the customer once all items are successfully installed. However, in the case that events 2, 3, or 4 occur then the mechanic will return the item to the stock room and issue another item, completing this cycle until event one happens and an item is successfully installed on the aircraft. Depending on whether or not the item is lost, scrapped or needs to be reworked it will either be found and returned to stock, be not found and thus, not returned to stock, be scrapped, or get sent off for rework and then returned to stock. Note that only in rare occasions is an item lost and not returned to stock. Also, note that the time between when an item is installed onto the aircraft until when the complete aircraft gets delivered to the customer it is considered work in progress. This leads us to our next discussion on an additional reason why this becomes such an interesting application of the problem.

When a shop floor mechanic goes to grab an item from the stock room and there is not one there to be issued, a shortage occurs, and the unfilled demand gets backordered. At this point in time

this delays the final delivery of the aircraft until this missing item is filled. It is not uncommon that a following manufacturing process may be reliant upon this missing item being installed. In this case, the aircraft is further delayed, as now not only is the final delivery of the aircraft delayed but the aircraft cannot move to the following work center for the next manufacturing process. When the aircraft is stopped on the manufacturing line due to a shortage, the aircraft may have thousands of parts already installed which are now stuck as work in progress. You can envision the case in which an aircraft is near completion and final delivery but then the shop floor mechanic goes to install a bracket with a nominal value but there is no bracket to be issued. This now almost complete aircraft becomes a very expensive piece of work in progress sitting out on the shop floor until another bracket can be procured. Hence, in our application of this problem it is not uncommon for an item to have a penalty cost associated with it that is magnitudes larger than the unit cost of the item. We will see in further sections that in these cases we obtain some interesting results from our mathematical models.

Another important aspect to consider about the aerospace supply chain is the structure of the purchase contracts involved. As discussed earlier, this results in the number of final products to be delivered relatively certain, resulting in a very low demand uncertainty for final products. However, raw materials, parts, and sub-assemblies will still be subject to scrap, loss, and rework. This drives variability in the model and result in demand uncertainty. Let us envision the following fictional example to further understand this concept. Say we need to deliver 5 additional XYZ aircrafts to complete a contract. The bill of materials (MRP schedule) says that we need 200 units of part A, 20 units of part B, 5 units of sub-assembly C, and 10 units of part D. However, all the parts will also have a unique probability distribution associated with them that relates to the

discrete random number of parts that is scrapped, lost, or reworked. The final quantity of each part that gets issued from the inventory stock room will be the number required per the MRP schedule plus a random quantity of parts that are scrapped, lost, or reworked which is generated from this true probability distribution. Generally, companies are aware of this phenomenon and track the historical number of parts that are scrapped, lost, or reworked. This example can be visualized in Table 2 below.

	MRP Schedule	Scrap, Loss, & Rework	Quantity Issued
Part A	200	2%	205
Part B	20	5%	21
Sub-assembly C	5	1%	5
Part D	10	0%	10

Table 2. Scrap, Loss & Rework Example

Here, we see that the quantity issued is not simply the MRP schedule quantity multiplied by the scrap, loss & rework percentage plus the MRP schedule quantity. To reiterate this important concept, this is because the percentage in Table 2 is based off of what has happened historically, while the actual number that gets scrapped, lost, or reworked is a function of the true random probability distribution relating to scrapped, lost or reworked parts. Attempts can be made to predict this variation and uncertainty in demand by applying distribution fitting algorithms to historical data. However, even with state-of-the-art distribution fitting software, it is impossible to 100% accurately model the true distribution due to the complex and random nature of processes that account for this variation.

The final thing to note about aerospace supply chains is that many of the parts and sub-assemblies for which we need to generate inventory policies for are also on contracts. According to a published news article, there are 40,000 procured parts actively on consumption-based ordering (CBO) contracts at The Boeing Company in the Saint Louis system at any given time (Cook, 2004). CBO contracts are designed to have the right number of parts and supplies on hand exactly when they are needed (Cook, 2004). CBO allows suppliers to view Boeing inventory systems and can view and monitor inventory levels whenever they choose, allowing a faster response times to fluctuating demands. The CBO system establishes a minimum and maximum level for which the on-hand inventory of the contracted part should stay between (Cook, 2004). It is the supplier's responsibility to ensure that the right number of parts are shipped at the right times to stay within this level (Cook, 2004). This also offers an advantage to the suppliers as they get to choose when and how many parts they want to ship (Cook, 2004). For instance, if parts are produced in batches on a CNC machine then the suppliers can plan such that only full batches are multiples of full batches are shipped at a time. This allows suppliers to maximize the utilization of their machining processes.

Chapter 2

Literature Review

This section will serve as the background research related to the mathematical inventory optimization models that are discussed in the methods section, along with an overview of other related concepts and ideas.

2.1 Newsvendor Model Research

The newsvendor model is a classical inventory model which has seen a plethora of exposure in research first dating back to 1888 when Edgeworth utilized the central limit theorem to calculate an optimal amount of cash reserves to satisfy random withdraws from depositors (Edgeworth 1888). The model was later coined the newsboy or newsvendor model to represent a newspaper vendor deciding how many copies of the paper to purchase given the uncertainty in demand. The newsvendor model shows some great promise in the problem definition presented in this thesis as we will be predicting a last period's inventory level given the uncertain demand due to scrap, loss, and rework. In the newsvendor model an optimal order quantity is calculated that should be ordered at the beginning of the period. A modern formulation of the model is seen in "Optimal inventory Policy" by (Arrow, Harris, & Marschak 1951). The original newsvendor model does not consider products that have a salvage value associated with them, which is often the case in the problem presented in this thesis. Fortunately, many efforts have been made to adjust the model to account for this. One particular formulation is obtained by including the salvage value in the underage cost (Nahmias & Olsen 2015). Another formulation for this is by adjusting the holding cost function to equal shortage cost minus salvage value (McGarvey, 2020).

2.2 (s, S) Inventory Policies Research

Seeing that aerospace companies use ordering policies similar to that of the CBO ordering policies discussed earlier, some (s, S) inventory policies are also explored. An (s, S) policy is one in which an order is placed whenever inventory reaches a level s with the quantity ordered bringing the inventory up to some level S (Zheng & Federgruen 1990). The s inventory level should always be less than that of the S inventory level for a plausible policy. Under an (s, S) policy, a set up cost and initial inventory are also considered, which was not the case in the newsvendor models.

This (s, S) policy is often referred to a minimum, maximum policy, which is the format that we see under the CBO and alike policies used across the aerospace industry. The (s, S) or Min/Max inventory policy has also been used extensively in other industries outside of the aerospace industry (McGarvey, 2019). An optimal (s, S) policy is found by minimizing the long-run average cost function (Veinott 1996). One formulation for finding optimal (s, S) policies is described by first finding S by the newsboy model formulation discussed earlier, then s is determined by solving a function related to the inventory policy costs (McGarvey 2019). A similar formulation is shown by Scarf (n.d.) An additional formulation using an algorithmic approach for finding optimal policies is defined (Zheng & Federgruen 1991). In this formulation, better bounds on optimality than in previous research are found. The better bounds significantly reduce the computation complexity, and the algorithmic approach presented claims a computation complexity of only 2.4 times that required to evaluate a single (s, S) policy (Zheng & Federgruen 1991). The algorithm applies to both periodic review and continuous review inventory systems (Zheng & Federgruen 1991).

2.3 Distribution Fitting

All the models considered here assume that demand is a random variable following a known or estimated probability distribution function. Finding a suitable distribution that fits the demand can sometimes be a challenging task. Traditionally, companies have used both Normal and Gamma distributions to fit production demand. Recently there has been a shift toward using the Gamma distribution for multiple different reasons. The first being that the Gamma distribution is very flexible and can mimic many other distributions such as the Normal and Poisson distributions which are commonly used to fit demand. In fact, the generalized Gamma distribution allows for the distribution to be shifted along the x axis by adding a third parameter. The Gamma distribution can also be convoluted to combine multiple distributions into a single distribution. This is particularly useful for the case when determining inventory levels that pool from multiple different products or departments.

2.3.1 Gamma Distribution

The Gamma distribution is a two-parameter continuous distribution commonly defined by either a shape and scale parameter or a shape and rate parameter. In this thesis the shape and rate parameter definition will be used. Using the Gamma distribution to fit demand is not a new concept and was in fact found to be a very good choice because of the flexibility and non-negative values that it takes on (Keaton 1995). The Gamma distribution is highly flexible and by adjusting the shape and rate parameters (and in some cases generalizing to allow for shifting), can assume practically any shape that can be expected by demand (Keaton 1995). The Normal distribution is typically a reasonable approximation for fast moving items however, not suitable for slow moving items as the distributions are typically skewed right (Keaton 1995). This is especially important in this analysis as aircrafts typically have a relatively slow manufacturing time, causing the inventory

levels managed to produce them to move slowly as well. The Gamma distribution is seen below for different levels of alpha or shape and beta or rate. Figure 7 shows the Gamma distribution with beta equal to one and Figure 8 with beta equal to 2.

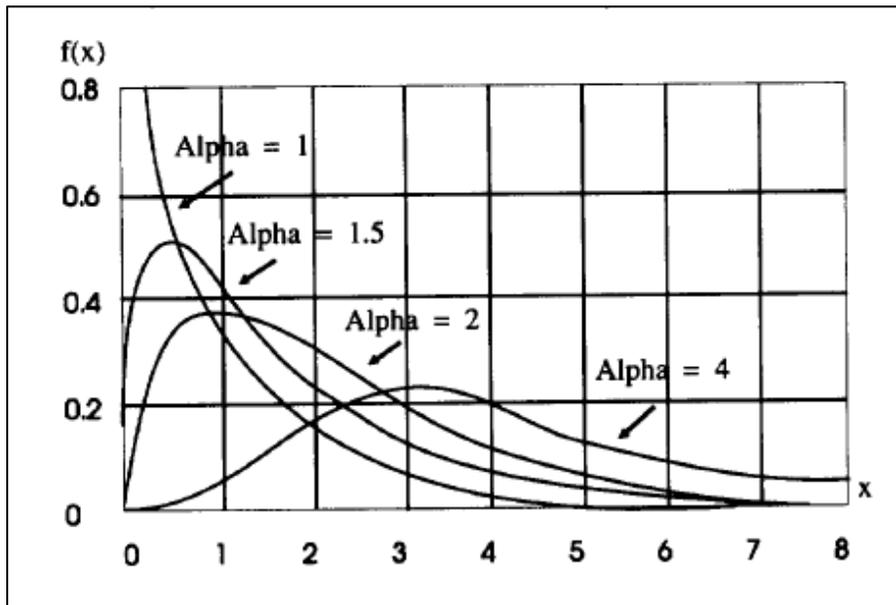


Figure 7. Gamma Density Function Beta = 1

Source: Mark Keaton, Gamma Density Functions; Beta = 1.

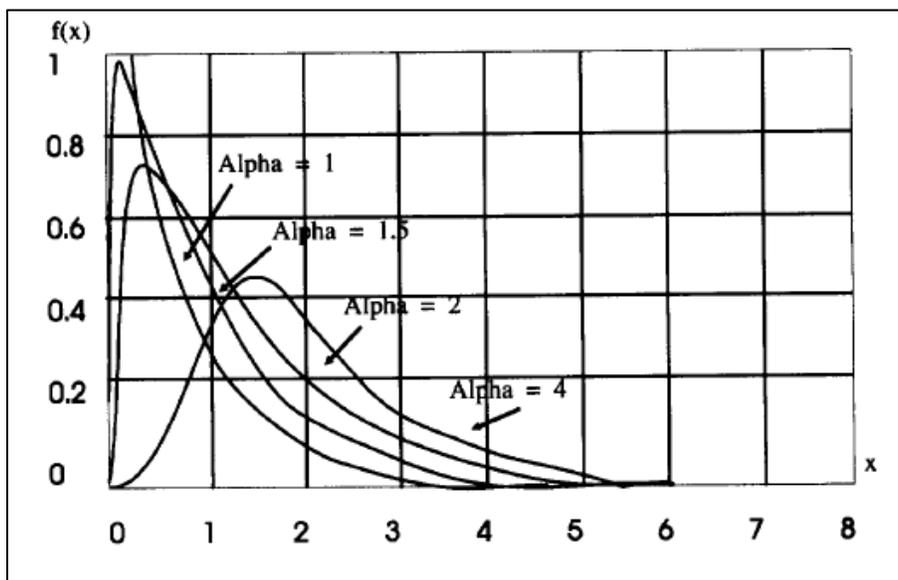


Figure 8. Gamma Density Function Beta = 2

Source: Mark Keaton, Gamma Density Functions; Beta = 2.

Another characteristic of the Gamma distribution that makes it a suitable choice for fitting demand is the relative ease in which one can calculate the service level measures with simple computer codes (Keaton 1995). Finally, the last reason why Gamma is a great choice for fitting demand is the ease for which one can estimate its parameters (Keaton 1995). A simple method in doing so is by first computing the sample mean and standard deviations represented by \bar{x} and s respectively (Keaton 1995). The method is represented by the equations below and termed the method of moments (Keaton 1995):

$$\alpha = \frac{\bar{x}^2}{s^2} \quad (1)$$

$$\beta = \frac{\bar{x}}{s^2} \quad (2)$$

However, this method does not give the most robust estimates and more computationally intense methods can be used to obtain the best fit (Keaton 1995). The method of moments is the least computational effort followed by “shortcut” methods of maximum likelihood with the most computationally intense being the full maximum likelihood method (Keaton 1995). Now we will visually demonstrate the flexibility of the Gamma distribution by applying the method of moments to three other distributions as seen in Keaton, 1995. Figure 9 shows this technique with the Poisson distribution, Figure 10 with the Binomial distribution, and Figure 11 with the shifted Lognormal and generalized (shifted) Gamma distribution (Keaton 1995).

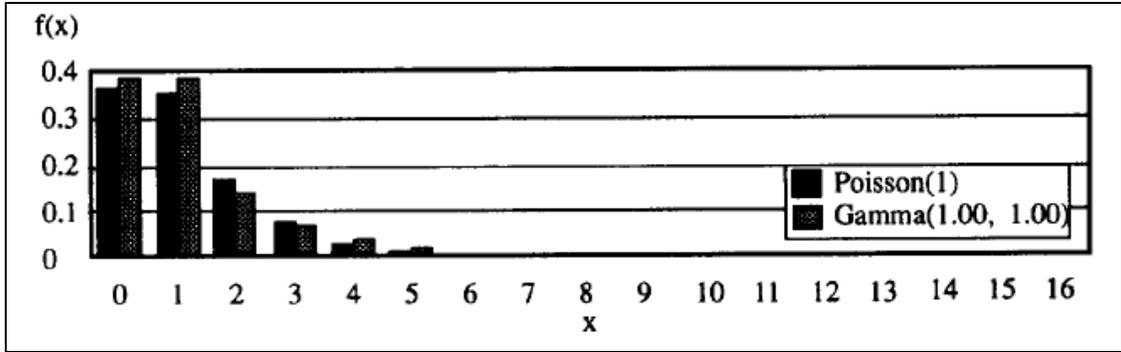


Figure 9. Comparison of Poisson and Gamma Distributions

Source: Mark Keaton, Comparison of Poisson and Gamma Distributions

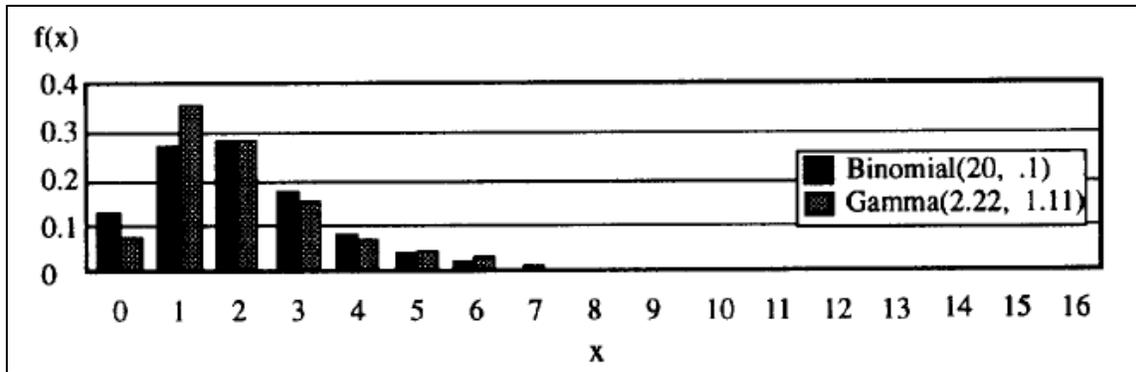


Figure 10. Comparison of Binomial and Gamma Distributions

Source: Mark Keaton, Comparison of Binomial and Gamma Distributions

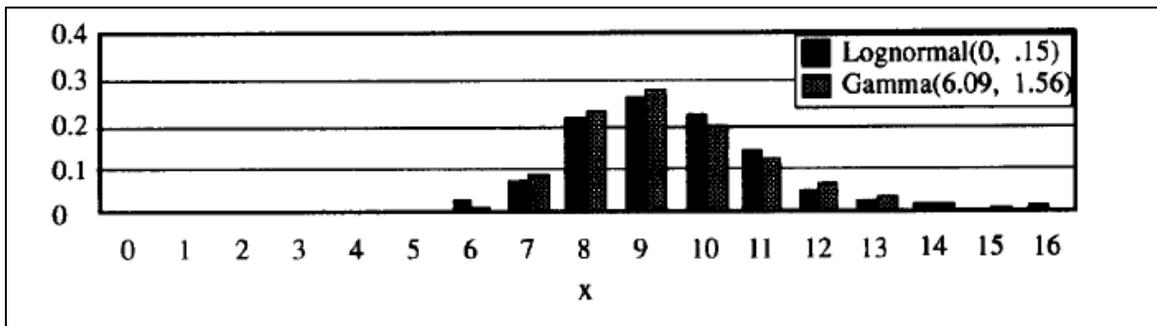


Figure 11. Comparison of Shifted Lognormal and Shifted Gamma Distributions

Source: Mark Keaton, Comparison of Shifted Lognormal and Shifted Gamma Distributions

As we can see, in all three instances above, the Gamma distribution provides a good estimate of the different distributions.

2.4 Shortages & Backorders

The inventory optimization models discussed earlier involve an estimation for the penalty cost associated with not having an item to issue at the time a demand needs to be satisfied. Often this penalty cost is very hard to estimate. This is because not only does the penalty cost represent the potential delay in the delivery of a final product, but it also represents the future loss in customer goodwill due to the delayed delivery (Liberopoulos et al, 2009). This loss in customer good can lead to a potential change in future demand which needs to be considered along with the current penalty associated with the delayed delivery (Schwartz 1966). In the case that the unsatisfied demand is backordered and filled at a late date, the direct and current cost relates to this delayed delivery and may also include extra administrative costs, material handling/transportation costs associated with the backordered demand being expedited, any contractual penalties, etc. (Liberopoulos et al, 2009). The direct costs have been shown to be successfully calculated with some effort, however, the indirect costs relating to loss of customer goodwill is much harder to estimate (Liberopoulos et al, 2009). Schwartz (1996) demonstrates a new model termed *perturbed demand* (PD) in which he modifies the economic order quantity model such that the penalty cost term is removed from the objective function and assumes that the long-run demand rate will be decreasing due to the loss in customer goodwill (Schwartz 1996). However, in the case that the product is nearing the end of a production run this approach will not be fitting, as there will be zero demand in following periods regardless of the loss in customer goodwill.

Chapter 3

Methodology

In this chapter, we will mathematically formulate the periodic review inventory models that are explored throughout this thesis. Parameters to the models are defined along with other supplemental information to the inventory models.

3.1 Parameter Definitions

In this section we will define important parameters that are used in the inventory model formulations that are explored. Note that the parameter definitions can be adjusted however the reader sees fit, but the mathematical models may also need to be adjusted to account for this.

3.1.1 Unit Ordering Cost (c)

Unit ordering cost (c) represents the cost incurred to procure one unit of an item. Here a fixed unit ordering cost will be used regardless of the quantity purchased.

3.1.2 Storage Cost

Storage cost represents the cost associated with all physical inventory on hand. The components of this cost include the cost of physical space required to store inventory, taxes and insurances on the inventory, breakage or deterioration of the inventory, opportuning costs of alternative investments, etc. (Nahmias & Olsen 2015). Commonly in practice storage cost will be an assigned percentage of the unit ordering costs by the finance department.

3.1.3 Salvage Value

Salvage value is the monetary amount that the company can expect to receive for an item in the case that there is excess inventory at the end of the period. This is particularly relevant in this problem as excess inventory cannot be used in a following periods demand. In the case that the item is used or sold by another department within the company, then the salvage value will be a percentage of the unit cost. The percentage represents the extra costs incurred by transferring the item to the different department. In the case that the item is not used or sold by a different department, then the salvage value is assumed to be zero, representing the previously purchased item being scrapped.

3.1.4 Holding Cost (h)

Holding cost (h) is defined as the cost per unit remaining at the end of the period. The holding cost will be equal to the storage cost minus salvage value. This is represented in the equation below:

$$\text{Holding Cost} = \text{Storage Cost} - \text{Salvage Value} \quad (3)$$

Note that in some cases when the item has a positive salvage value, the holding cost may become negative. This will be an interesting case examined later.

3.1.5 Penalty Cost (p)

Penalty cost represents the shortage costs discussed in the literature review. Recall that this value represents both the direct cost associated with not having inventory on hand to satisfy demand and the indirect cost in loss of customer goodwill. Penalty cost is defined here as the cost associated

with not having sufficient stock on hand to satisfy demand (Nahmias & Olsen 2015). In this analysis it will be a fixed value assigned to each item analyzed.

3.1.6 Set Up Cost (k)

Set up cost (k) refers to the costs associated with ordering the item, this cost generally represents things such as administrative costs incurred placing the order, shipping, and handling, delivering of the product, etc. In this analysis a fixed set up cost will be assigned to each item.

3.1.7 Initial Inventory (x)

Because we are examining periodic review inventory models, initial inventory (x) represents the quantity of inventory on hand at the point in time considered.

3.2 Lead Time Scaling

Often demand is forecasted on a periodic basis which could be days, month, years, etc. Due to this there is a need to scale the demand distribution to fit lead time. In the case that demand is assumed to follow a Normal distribution, the sums of independent random variables are also normally distributed and thus, the lead time distribution is normal (Nahmias & Olsen 2015). For this case all we have to do is determine the distribution parameters. Because both the means and variances are additive, the formulas for scaling to lead time are those shown below with sample mean (\bar{x}) and sample standard deviation (s):

$$\bar{x}_{lead\ time} = \bar{x}_{period} * (lead\ time\ in\ periods) \quad (4)$$

$$s_{lead\ time} = s_{period}\sqrt{lead\ time\ in\ periods} \quad (5)$$

In the case that we use the Gamma distribution to fit demand, we simply scale to lead time using the Normal distribution and then apply the method of moments discussed in section 2.3.1.

3.3 Newsvendor Model Formulation

In this section a newsvendor model formulation tailored to the specific problem at hand will be presented. The formulation presented here relies on the formulation seen in McGarvey (2019) and Nahmias and Olsen (2015) with adjustments to include an initial inventory level. The period analyzed is one lead time away from the end of production and thus, periodic mean and standard deviation values for demand are scaled as discussed in section 3.2.

Notation:

- c = Unit ordering cost.
- h = Holding cost.
- p = Penalty cost.
- x = Inventory on hand at the beginning of the period.
- D = Random demand occurring in the analyzed period
- Probability distribution of demand $P_D(D)$.

Assumptions of the model:

- Demand in this period is a random variable with a known probability distribution.
- Holding and shortage costs are linear

Decision variable:

- The quantity of units to have on hand at the beginning of the period represented by y .

Objective Function:

- Minimize the expected total cost function represented by $E[TC(y)]$.

Constraints of the model:

- Decision variable y must always be greater than or equal to $(y \geq x)$.

It is important to understand that because demand is a random variable, total cost is also a random variable. If demand is a discrete random variable with cumulative distribution function of the form shown below:

$$F_D = \sum_{d=0}^y P_d(D) \quad (6)$$

Then expected total cost is:

$$\begin{aligned} E[TC(y)] &= \sum_{d=0}^{\infty} [(c * (y - x) + p * \max\{0, d - y\} + h * \max\{0, y - d\})P_D(d)] \\ &= c * (y - x) + \sum_{d=y}^{\infty} (p(d - y)P_D(d)) + \sum_{d=0}^{y-1} (h(y - d)P_D(d)) \\ &= c * (y - x) + L(y) \end{aligned} \quad (7)$$

With optimal inventory to have on hand at the beginning of the period y^* being the smallest integer that satisfies the following equation and the constraint listed above:

$$F_D(y^*) \geq \frac{p - c}{p + h} \quad (8)$$

In the case that the above equation produces a value less than the initial inventory level, simply set y^* equal to the initial inventory level. That is, do not procure any additional units of inventory as you are already at an optimal solution given the initial inventory level.

For the case where we assume demand follows a continuous distribution such as the Normal or Gamma distributions, we can extend the above formulation and approximate the discrete demand by a continuous random variable D , with probability density function $f_D(z)$. The cumulative distribution of D is shown in the equation below.

$$\Phi(z) = \int_0^z f_D(y) dy \quad (9)$$

The expected total cost then becomes:

$$\begin{aligned} E[TC(y)] &= \int_0^\infty [(c * (y - x) + p * \max\{0, z - y\} + h \\ &\quad * \max\{0, y - z\})f_D(z)]dz \\ &= c * (y - x) + \int_y^\infty (p(z - y)f_D(z))dz + \int_0^y (h(y - z)f_D(z))dz \\ &= c * (y - x) + L(y) \end{aligned} \quad (10)$$

With optimal inventory to have on hand at the beginning of the period y^* satisfying the following equation and the constraint listed above:

$$\Phi(y^*) = \frac{p - c}{p + h} \quad (11)$$

In the case of the continuous random variable, the y^* value calculated is rarely of integer form. If the optimal y^* is not an integer, simply round up to the nearest integer for a plausible inventory level, hence, the approximation using continuous random variables for demand. The term $L(y)$ in the expected total cost functions represents the expected total shortage plus holding costs. While the $\frac{p-c}{p+h}$ term represents the critical ratio or optimal service level to meet.

Given the solution y^* , we can now calculate the quantity of units to be procured to ensure that we arrive to an optimal inventory level at the beginning of the period represented by λ . This is done by simply subtracting the optimally calculated value of units to have on hand at the beginning of the period by the current level of inventory on hand. This is shown below:

$$\lambda = y^* - x \quad (12)$$

3.4 (s, S) Model Formulations

In this section both of the (s, S) models discussed previously will be explored, starting with the equation-based method first followed by the algorithm based method.

3.4.1 (s, S) Equation Based

Here we will formulate the (s, S) inventory model based on the equation-based model discussed in section 2.2. Similarly, to the newsvendor model, this formulation will rely on that seen in

McGarvey (2019) and Scarf (n.d.). Essentially this is an extension of the newsvendor model formulated earlier with the additional consideration of a set up cost K . Lead time periods are used in this model so periodic mean and standard deviation values for demand are again scaled as discussed in section 3.2.

Notation:

- c = Unit ordering cost.
- h = Holding cost.
- p = Penalty cost.
- K = Set up cost.
- x = Initial inventory on hand.
- D = Random demand occurring within a period D
- Probability distribution of demand $P_D(D)$.

Assumptions of the model:

- Demand is a random variable with a known probability distribution.
- Holding and shortage costs are linear.

Decision variables:

- Inventory levels s and S

The same $L(y)$ term used to represent the expected holding and shortage costs in the newsvendor model is again used here. In the case of discrete random demand:

$$L(y) = \sum_{d=y}^{\infty} (p(d - y)P_D(d)) + \sum_{d=0}^{y-1} (h(y - d)P_D(d)) \quad (13)$$

In the case of continuous random demand:

$$L(y) = \int_y^{\infty} (p(z - y)f_D(z))dz + \int_0^y (h(y - z)f_D(z))dz \quad (14)$$

The total expected total inventory cost incurred by bringing inventory up to level y from x is:

$$K + c(y - x) + L(y) \quad \text{if } y > x \quad (15)$$

$$L(y) \quad \text{if } y = x \quad (16)$$

Here we will define s and S as follows:

- S is then the value of y that minimizes $cy + L(y)$
- While s is the smallest value of y for which $cs + L(s) = K + cS + L(S)$

Now we will investigate how the optimal inventory policy is determined:

We can see that if $x > S$ then:

$$K + cy + L(y) > cx + L(x) \text{ for all } y \geq x \quad (17)$$

This can be rearranged to the form of:

$$K + c(y - x) + L(y) > L(x) \quad (18)$$

Note that the left-hand side of this inequality is representative of the expected total inventory cost of ordering $y - x$ units which brings the inventory level up to the quantity y . While the right-hand side represents the expected total inventory cost if no additional inventory is ordered. Hence, the optimal policy is one in which if $x > S$, do not order any additional units of inventory.

Similarly, if $s \leq x \leq S$ then:

$$K + cy + L(y) \geq cx + L(x) \text{ for all } y > x \quad (19)$$

Again, this is rearranged to the form:

$$K + c(y - x) + L(y) \geq L(x) \quad (20)$$

We see the same results as before, leading us to believe that the optimal policy will not order any additional inventory when $s \leq x \leq S$.

Lastly, if $x < s$ then:

$$\min_{y \geq x} \{K + cy + L(y)\} = K + cS + L(S) < cx + L(x) \quad (21)$$

Which can be rearranged to:

$$\min_{y \geq x} \{K + c(y - x) + L(y)\} = K + c * (S - x) + L(S) < L(x) \quad (22)$$

This shows us that minimum cost is incurred by bringing the inventory level up to S .

The optimal inventory policy is summarized below (McGarvey 2019):

- If $x < s$, order $S - x$ additional units
- If $x \geq s$, do not order any additional units of inventory.

S is then obtained by:

$$\Phi(S) = \frac{p - c}{p + h} \quad (23)$$

s is obtained by finding the smallest value for s that satisfies:

$$cs + L(s) = K + cS + L(S) \quad (24)$$

3.4.1.1 (s, S) Equation Based with Normal Distribution

By assuming that demand follows a Normal distribution, we will need to use the Normal distributions probability density function in the definition of $L(y)$. To do this we simply replace the $f_d(z)$ term with the probability distribution function of the normal distribution. The probability distribution function of the Normal distribution is seen below:

$$f(z) = \frac{1}{\sigma\sqrt{2\pi}} * e^{-\frac{1}{2}\left(\frac{z-\mu}{\sigma}\right)^2} \quad (25)$$

Substituting this value in for $f_d(z)$ we obtain $L(y)$ as follows:

$$\begin{aligned} L(y) = & \int_y^\infty p(z-y) \frac{1}{\sigma\sqrt{2\pi}} * e^{-\frac{1}{2}\left(\frac{z-\mu}{\sigma}\right)^2} dz \\ & + \int_0^y h(s-z) \frac{1}{\sigma\sqrt{2\pi}} * e^{-\frac{1}{2}\left(\frac{z-\mu}{\sigma}\right)^2} dz \end{aligned} \quad (26)$$

We then obtain an adjusted equation that solves for s by substituting the above function for $L(y)$:

$$\begin{aligned} & cs + \int_s^\infty p(z-s) \frac{1}{\sigma\sqrt{2\pi}} * e^{-\frac{1}{2}\left(\frac{z-\mu}{\sigma}\right)^2} dz \\ & \quad + \int_0^s h(s-z) \frac{1}{\sigma\sqrt{2\pi}} * e^{-\frac{1}{2}\left(\frac{z-\mu}{\sigma}\right)^2} dz \\ = & K + cS + \int_s^\infty p(z-S) \frac{1}{\sigma\sqrt{2\pi}} \\ & \quad * e^{-\frac{1}{2}\left(\frac{z-\mu}{\sigma}\right)^2} dz \int_0^S h(S-z) \frac{1}{\sigma\sqrt{2\pi}} * e^{-\frac{1}{2}\left(\frac{z-\mu}{\sigma}\right)^2} dz \end{aligned} \quad (27)$$

3.4.1.2 (s, S) Equation Based with Gamma Distribution

Similarly, we can do the same with the Gamma distribution for assuming the demand follows the Gamma distribution with shape (α) and rate (β) parameters. The probability distribution for the Gamma distribution is seen below:

$$f(z) = \frac{\beta^\alpha}{\Gamma(\alpha)} z^{\alpha-1} e^{-\beta z} \quad (28)$$

With $\Gamma(\alpha)$ defined as:

$$\Gamma(\alpha) = \int_0^\infty z^{\alpha-1} e^{-z} dz \quad (29)$$

We then obtain $L(y)$ as:

$$\begin{aligned} L(y) &= \int_y^\infty p(z-y) * \frac{\beta^\alpha}{\Gamma(\alpha)} z^{\alpha-1} e^{-\beta z} dz \\ &+ \int_0^y h(s-z) \frac{1}{\sigma\sqrt{2\pi}} * \frac{\beta^\alpha}{\Gamma(\alpha)} z^{\alpha-1} e^{-\beta z} dz \end{aligned} \quad (30)$$

Again, we obtain are adjusted equation that solves for s by substituting the above function for

$L(y)$:

$$\begin{aligned} cs + \int_s^\infty p(z-s) \frac{\beta^\alpha}{\Gamma(\alpha)} z^{\alpha-1} e^{-\beta z} dz \\ + \int_0^s h(s-z) \frac{\beta^\alpha}{\Gamma(\alpha)} z^{\alpha-1} e^{-\beta z} dz \\ = K + cS \\ + \int_s^\infty p(z-S) \frac{\beta^\alpha}{\Gamma(\alpha)} z^{\alpha-1} e^{-\beta z} dz \int_0^S h(S-z) \frac{\beta^\alpha}{\Gamma(\alpha)} z^{\alpha-1} e^{-\beta z} dz \end{aligned} \quad (31)$$

3.4.2 (s, S) Algorithm Based

Here we will define the algorithm used to find optimal solutions for an (s, S) inventory policy mentioned in section 2.2. The algorithm formulation is originally the same formulation as used in Zheng and Federgruen (1991) with a slightly different notation.

Notation:

- D = The one random variable corresponding to the one-period demand
- $p_j = \Pr\{D = j\}$, $j = 0, 1, 2, \dots$;
- K = Set up cost.
- $L(y)$ = The one-period expected holding and shortages costs when starting with inventory position y (This is the same function as used in previous formulations).
- c = Unit ordering costs

Assumptions of the model:

- All stockouts are backordered.
- One-period demands are independent and independently distributed and of integer value.
- Holding, shortage, and unit ordering costs are stationary and increase linearly or convexly with the end of period shortage size.
- That $-L(\cdot)$ is unimodal (Thus, $L(\cdot)$ is convex).
- $\lim_{|y| \rightarrow \infty} L(y) > \min_y L(y) + K$

Decision variables:

- Inventory levels s and S .

Objective Function:

- Minimize the long-run average costs, $c(s, S)$.

The long-run average cost function $c(s, S)$ is defined as:

$$c(s, S) = M(S - s)K + \sum_{j=0}^{S-s-1} m(j)L(S - j) \quad (32)$$

With $m()$ defined as:

$$m(0) = (1 - p_0)^{-1} \quad (33)$$

$$m(j) = \sum_{l=0}^j p_l m(j - 1), \quad j = 1, 2, \dots \quad (34)$$

And $M()$ defined as:

$$M(0) = 0 \quad (35)$$

$$M(j) = M(j - 1) + m(j - 1), \quad j = 1, 2, \dots \quad (36)$$

The Algorithm:

Step 0. $s := y^*$;

$S_0 := y^*$;

Repeat $s := s - 1$ until $c(s, S_0) \leq L(s)$;

$s_0 := s$; $c^0 := c(s_0, S_0)$; $S^0 := S_0$; $S := S^0 + 1$;

Step 1. While $L(S) \leq c^0$ do

begin If $c(s, S) < c^0$

then begin $S^0 := S$.

```

    While  $c(s, S^0) \leq L(s + 1)$  do  $s := s + 1$ ;
     $c^0 = c(s, S^0)$ ;
    end;
     $S := S + 1$ ;
end.

```

In step 0 of the algorithm, y^* is simply the solution of a one-period newsboy problem (Zheng & Federgruen 1991). In step 0 an initial order-up-to level is entered, and then an optimal corresponding reorder level is found by incrementally decreasing by one (Zheng & Federgruen 1991). In step 1 a better value of S is searched for by incrementing by one and if a better value is found, the new value is updated and a new optimal reorder level is found by incrementing the old value by one (Zheng & Federgruen 1991). In the last iteration of the algorithm, $S^0 = S^*$ and $s^0 = s^*$ for an optimal policy (s^*, S^*) with $c^0 = c^*$ (Zheng & Federgruen 1991). For more information, pertaining to the algorithm, refer to Zheng and Federgruen (1991).

3.4.2.1 (s, S) Algorithm Based Implementation Challenges

In this section we discuss some of the challenges faced while trying to implement the algorithm and some steps we took to tailor the algorithm to the specific problem at hand. The first thing to note is that, this algorithm was not developed to consider a positive salvage value associated with the item being analyzed. This is especially true for the case when the salvage value is great enough to cause the holding cost to become negative. An interesting situation was encountered while running the algorithm on a part that meets this description. The algorithm continuously increased the inventory levels indefinitely. This happened because of the invalid assumptions that assume that $-L(\cdot)$ is unimodal and thus, $L(\cdot)$ is convex and $\lim_{|y| \rightarrow \infty} L(y) > \min_y L(y) + K$. At some point

in time when inventory is increased to an amount where the possibility of a shortage converges to zero, the only additional cost being added to $L(y)$ is the holding cost. In the case that the holding cost is negative, this is driving our $L(y)$ to start decreasing as y increases past this point where the probability of a shortage is essentially zero. This continues, eventually driving the whole term $L(y)$ to become negative. If we take another look at the assumption regarding limits, this is obviously breaking that assumption because as y approaches infinity, $L(y)$ is strictly decreasing. Now, if we consider the assumptions regarding the curvature of $L(y)$ we see that it is assumed that $L(y)$ is a convex function. Let's take a look at some graphs regarding the expected holding cost, shortage cost, and holding plus shortage cost or $L(y)$ of a part that has a salvage value such that the holding cost becomes negative to further understand this concept. Figure 12 shows the graph of the expected holding costs, Figure 13 of the expected shortage costs, and Figure 14 of $L(y)$.

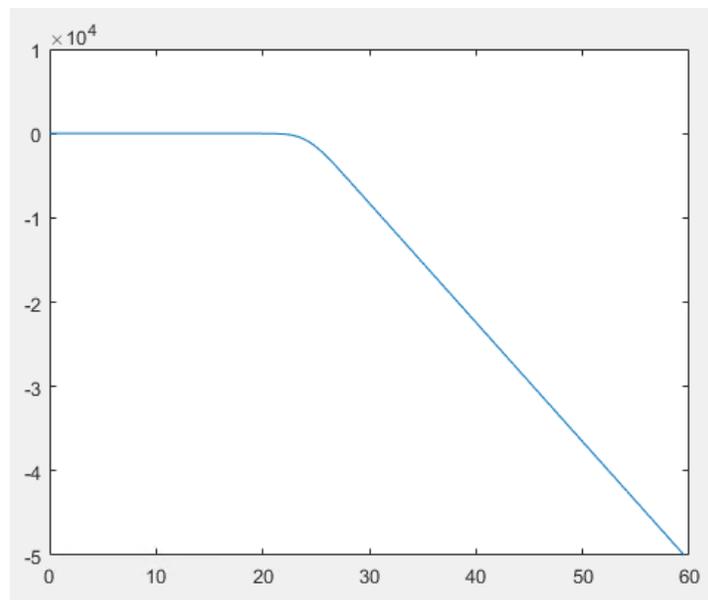


Figure 12. Graph of One Period Expected Holding Cost Vs. Inventory Position

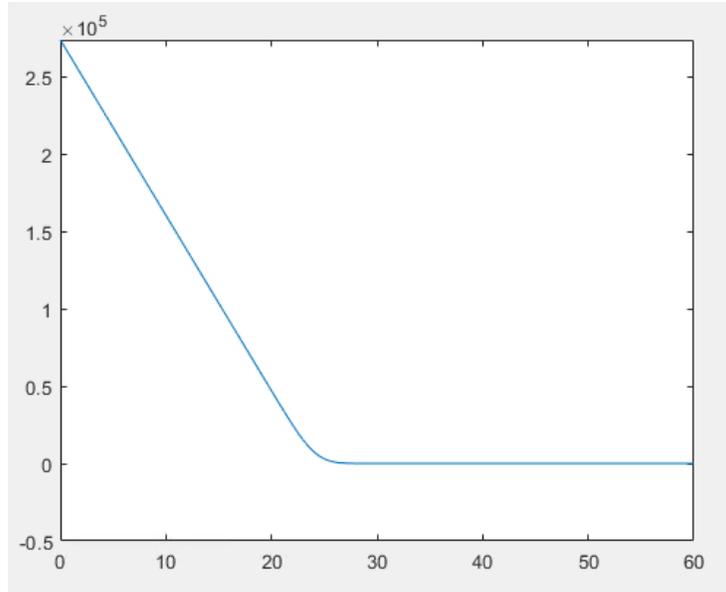


Figure 13. Graph of One Period Expected Shortage Cost Vs. Inventory Position

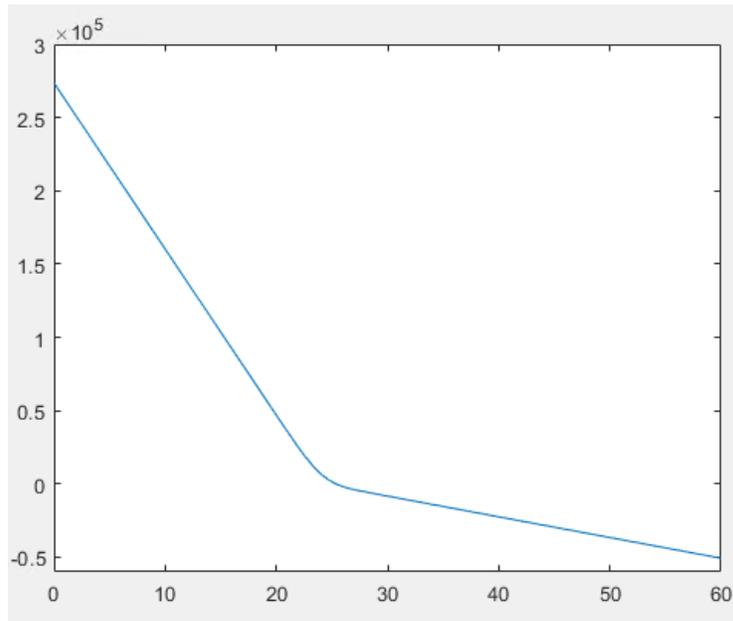


Figure 14. Graph of $L(y)$ Vs. Inventory Position

As we can see from the graphs above, for this particular case, the shortage costs probability seems to converge to zero around 25 units of inventory, from this point on both the expected holding cost and $L(y)$ start to become strictly decreasing, causing both the assumptions mentioned to be broke. According to Zheng and Federgruen (1991) “The long-run average order quantity equals ED under any policy that avoids infinitely large inventories or backlogs. Linear order costs may thus be ignored for the purpose of determining an optimal policy” (p. 655). Infinitely large inventories are the case here, so the unit ordering cost is added to the model to see if this fixes the issue. This is done by simply including the unit ordering costs into the $L(y)$ function by:

$$G(y) = cy + \int_y^{\infty} (p(z - y)f_D(z))dz + \int_0^y (h(y - z)f_D(z))dz \quad (37)$$

Here $G(y)$ now represents the one period expected holding plus shortage plus unit ordering costs. Now we look at the new graph produced with the same data values as that used previously in Figure 15.

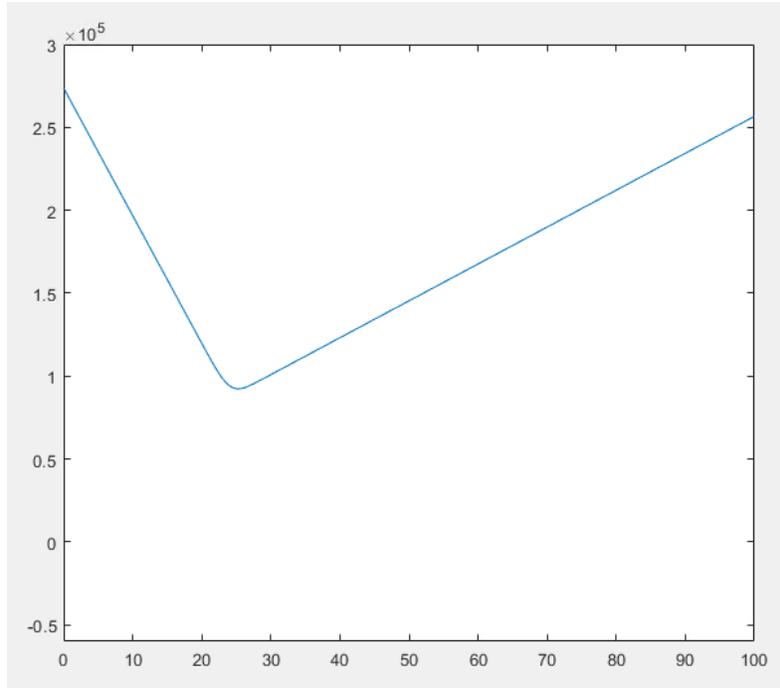


Figure 15. Graph of $G(y)$ Vs. Inventory Position

It appears as if both assumptions that were previously broken are now valid. However, while running the algorithm we run into an additional issue created by our redefinition of the one period expected costs. The algorithm now gets stuck for long periods in the outer while loop of step one. This is particularly true for parts with small unit ordering costs. Refer to Figure 16, where the orange line represents c^0 and the blue line represents $G(S)$.

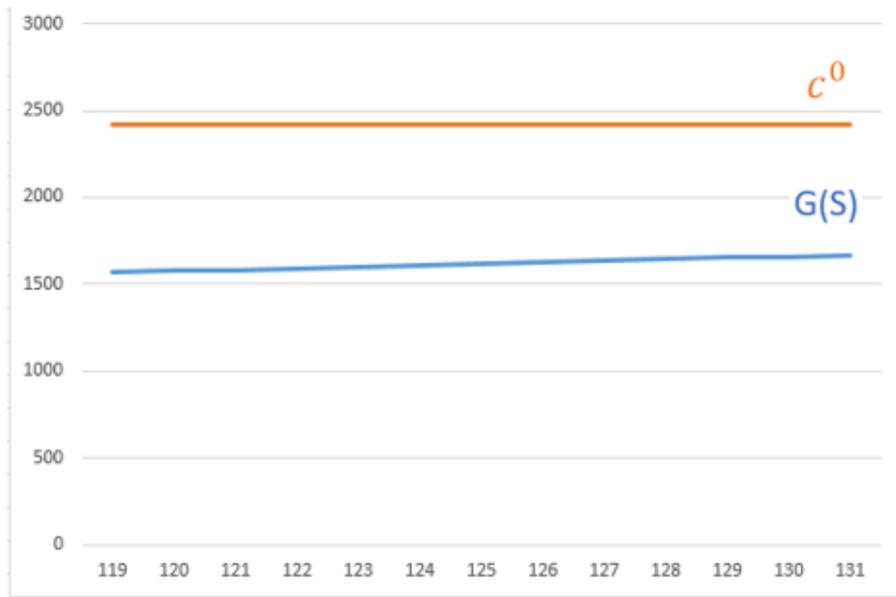


Figure 16. Graph of $G(S)$ and c^0 Vs. Inventory Position

A low-cost part is used to produce the graph. As we see, the c^0 value never changes because the cost function never gets updated. The slope of $G(S)$ is relatively low causing the number of iterations to happen until the two curves converge and the while loop breaks to take a significant amount of computation time. I have found a solution to this; however, it changes the definition of the algorithm to work as a heuristic instead of guaranteeing to always find an optimal solution. The solution proposed is to cap the number of iterations that occur in this while loop to some number of standard deviations away from the mean. This change is shown in the new algorithm formulation, using both the new $G(y)$ function and additional break if statement. The statement is bolded to allow for easy reference to the reader. Where standard deviation represents the lead, time adjusted standard deviation for the given item being analyzed and φ represents the number of standard deviations to cap to. A higher number will result in higher confidence that the produced solution is optimal but also take more computation time, while a lower value lowers the chance that the found solution is optimal but decreases the computation time.

The Algorithm:

Step 0. $s := y^*$;

$S_0 := y^*$;

Repeat $s := s - 1$ *until* $c(s, S_0) \leq G(s)$;

$s_0 := s$; $c^0 := c(s_0, S_0)$; $S^0 := S_0$; $S := S^0 + 1$;

Step 1. *While* $G(S) \leq c^0$ *do*

begin if $c(s, S) < c^0$

then begin $S^0 := S$.

While $c(s, S^0) \leq L(s + 1)$ *do* $s := s + 1$;

$c^0 = c(s, S^0)$;

end;

$S := S + 1$;

break if $S - y^* > \varphi * (\textit{standard deviation})$

end.

end.

A final thing to note about the algorithm-based method is that the Zheng and Federgruen (1991) provide 24 different example problems and the solutions to each. The first example problem will be referenced to show this discrepancy. These are shown in Table 3 below.

μ	s^*	S^*	c^*	y^*	\underline{s}	s_0	\bar{S}	\bar{S}^*	$\bar{S} - \underline{s}$	u	IS	B	# +	# *	# /	#COM	#TOT	#SNG	\hat{R}
10	6	40	35.022	14	3	3	53	45	50	39	26	3,650	1,559	1,564	45	129	3,297	1,892	1.74
15	10	49	42.698	20	7	7	83	57	76	47	29	5,118	2,240	2,247	53	149	4,689	2,652	1.77
20	14	62	49.173	26	12	12	92	69	80	55	29	6,626	3,031	3,041	59	161	6,292	3,422	1.84
25	19	56	54.262	32	16	16	98	79	82	60	19	8,041	3,621	3,631	66	152	7,470	4,160	1.80
30	23	66	57.819	37	21	21	103	87	82	64	16	8,830	4,098	4,110	68	151	8,427	4,556	1.85
35	28	77	61.215	43	26	26	109	96	83	68	15	9,907	4,626	4,639	72	156	9,493	5,112	1.86
40	33	87	64.512	48	31	31	115	104	84	71	13	10,783	5,061	5,074	75	158	10,368	5,550	1.87
45	37	97	67.776	54	36	36	121	112	85	75	12	11,655	5,584	5,600	77	160	11,421	6,006	1.90
50	42	108	70.975	59	41	41	126	120	85	78	10	12,603	6,053	6,069	80	162	12,364	6,480	1.91
55	47	118	74.149	65	46	46	132	129	86	82	8	13,884	6,686	6,703	84	165	13,638	7,140	1.91
60	52	129	77.306	70	51	51	137	137	86	85	4	14,916	7,200	7,217	87	163	14,667	7,656	1.92
65	56	75	78.518	75	56	56	143	143	87	87	0	15,265	7,467	7,486	87	156	15,196	8,732	1.94
70	62	81	79.037	81	62	62	149	149	87	87	0	15,265	7,467	7,486	87	156	15,196	7,832	1.94
75	67	86	79.554	86	67	67	154	154	87	87	0	15,265	7,467	7,486	87	156	15,196	7,832	1.94
21	15	65	50.406	27	13	13	93	71	80	56	27	6,862	3,145	3,155	60	159	6,519	3,540	1.84
22	16	68	51.632	28	14	14	94	73	80	57	26	7,102	3,262	3,272	61	159	6,754	3,660	1.85
23	17	52	52.757	29	15	15	95	75	80	58	23	7,346	3,380	3,390	62	155	6,987	3,782	1.85
24	18	54	53.518	30	15	15	96	77	81	59	21	7,818	3,513	3,522	65	155	7,255	4,032	1.80
51	43	110	71.611	60	42	42	127	122	85	79	10	12,927	6,213	6,229	81	164	12,687	6,642	1.91
52	44	112	72.246	61	43	43	129	124	86	80	9	13,255	6,374	6,390	82	164	13,010	6,806	1.91
59	51	126	76.679	69	50	50	136	135	86	84	5	14,568	7,027	7,044	86	163	14,320	7,482	1.91
61	52	131	77.929	71	52	52	139	138	87	86	2	14,913	8,293	7,312	86	158	14,849	7,656	1.94
63	54	73	78.287	73	54	54	141	141	87	87	0	15,265	7,467	7,486	87	156	15,196	7,832	1.94
64	55	74	78.402	74	55	55	142	142	87	87	0	15,265	7,467	7,486	87	156	15,196	7,832	1.94

u = the maximum index j for which $m(j)$ is computed;
 IS = the number of improving S values;
 B = the bound for the number of operations needed as defined in the text;
+ (*, /) = the actual number of additions (multiplications, divisions);
#COM = the actual number of comparisons;
#TOT = the actual total number of elementary operations (= (# +) + (# *) + (# /) + (#COM));
#SNG = the number of operations needed for computing a single cost value $c(s_0, \bar{S}^*)$;
 \hat{R} = #TOT/#SNG.

Table 3. (s, S) Algorithm Based Example Problems

Source: A. Federgruen and Yu-Sheng Zheng, Performance of Algorithm on 24 Test Problems

In all example problems, it is stated that linear holding and shortage costs are used, lead time is zero, and Poisson distributed one-period demands (Zheng & Federgruen 1991). In all 24 test problems, a fixed set up cost of 24, holding cost rate of 1, and penalty cost rate of 9 are used (Zheng & Federgruen 1991). Thus, the problems with respect to the mean one-period demand represented by μ in the table (Zheng & Federgruen 1991). We see that in the first problem, $\mu = 10$, $s^* = 6$, $S^* = 40$, and $c^* = 35.022$. However, when I calculated the cost function $c^* = c(s^*, S^*)$ my calculated answer is much larger than the value shown in the table, in fact I

calculate cost to be $c^* = 166.27$. I show the cost function $c = c(s, S)$ again below to reference as the discrepancy is investigated:

$$c(s, S) = M(S - s)K + \sum_{j=0}^{S-s-1} m(j)L(S - j) \quad (38)$$

With $m()$ defined as:

$$m(0) = (1 - p_0)^{-1} \quad (39)$$

$$m(j) = \sum_{l=0}^j p_l m(j - 1), \quad j = 1, 2, \dots \quad (40)$$

And $M()$ defined as:

$$M(0) = 0 \quad (41)$$

$$M(j) = M(j - 1) + m(j - 1), \quad j = 1, 2, \dots \quad (42)$$

Here, the standard definition of $L(y)$ is used and no order costs are added to the model. This is simply the formulation as stated by Zheng and Federgruen (1991). We see that the cost function value must be $c(s, S) \geq M(S - s)K$ as the $\sum_{j=0}^{S-s-1} m(j)L(S - j)$ term should never be negative given the formulation. In the case that $s = 6$, $S = 40$, and $K = 24$ then, $c(s, S) \geq M(34) * 24$. Noting that, both $m(j)$ and $M(j)$ are strictly increasing by definition, in the case that $\mu = 10$, $m(0) \approx 1$ and increases from there with $M(34) \approx 3.9002$. Thus, $M(34) * K \approx 3.9002 * 24 \approx$

93.6 which implies that $c(6, 40) \geq 93.6$. This is all done just by definition of the cost function $c(s, S)$, $m(j)$, and $M(j)$ without relying on any distributions or algorithms to cause potential errors in calculations. It is for this reason that this method was ruled out for use in actual analysis, however, good solutions were still obtained given the redefinition of the algorithm which will be presented in the results section for example purposes only.

3.5 Computational Methods

In this section we discuss the methods for which the mathematical calculations were made. It was desired to complete as much as the computation as possible using Microsoft Excel software. This is because Microsoft Office is used extensively throughout industry and academia and thus, it offers the advantage of being a familiar for a large population. However, because of the computation required to compute some of the functions, such as those in the (s, S) inventory policies, it was decided that MATLAB software would be used along with Excel. This significantly simplifies the process and computation speed compared to that of using Excel with VBA code. The newsvendor model can easily be computed entirely within Excel using built in functions such as NORM.INV and GAMMA.INV. However, to compute the integrals and solve the equations defined for the (s, S) policies, MATLAB software was used. All simple calculations such as computing the holding cost, critical ratio, etc. were done in Excel and then the values were inputted into MATLAB for more complex analysis of the integrals involved. The complete algorithm mentioned in the previous section was developed using MATLAB, as it relies heavily upon the more complex functions that must be recalculated at several steps throughout the algorithm. Figure 17 below shows a screenshot of one of the equation-based method scrips to give the reader a visualization of the MATLAB interface and code structure.

```

18
19 %Declare number of calculations to be made
20 num_calc = numel(max);
21
22 %Loop through calculations
23 for i = 1:num_calc
24
25
26 %If max is zero (Covers the zero demand case or case when Max
27 %Calculated as Zero)
28 if max(i) == 0
29     min(i) = 0;
30
31 else
32 %Define pdf dist for normal evaluated at x
33     dist1 = pdf('Normal',x,mean(i),std(i));
34
35 %Define Expected Shortage Cost Function (Pre Integral)
36     s_cost = inline(penalty_cost(i)*(x-max(i))*dist1);
37     s_cost1 = inline(penalty_cost(i)*(x-s)*dist1, 's', 'x');
38
39 %Define Expected Holding Cost Function (Pre Integral)
40     h_cost = inline(holding_cost(i)*(max(i)-x)*dist1);
41     h_cost1 = inline(holding_cost(i)*(s - x)*dist1, 's', 'x');
42
43 %set eqn to be cs + L(s) = K + cS + L(S)
44     eqn = ordering_cost(i)*s + int(s_cost1(s, x),x,s,inf) + int(h_cost1(s, x),x,0,s) == set_
45
46 %solve equation numerically
47     min(i) = vpasolve(eqn, s);
48 end
49 end
50
51 %convert solution to double variable from a sym variable
52 min1 = double(min)
53

```

Figure 17. MATLAB Script Example

Note that some of the parameters used in Figure 16 may be different than that in the formal formulations provided in previous sections.

Chapter 4

Results & Analysis

4.1 Data

All data shown throughout this thesis is fictional and for example purposes only. With that being said, the example dataset shown below in Table 4 is used in example problems that demonstrate the mathematical models explored.

Item #	Salvage?	Unit Cost	Inventory Level	Mean	Std	PLT (days)
1	No	\$3.23	4	129	6.25	87
2	Yes	\$96,444.00	0	7	0.52	648
3	No	\$500.49	0	97	4.57	212
4	Yes	\$28,209.00	45	43	3.67	342
5	Yes	\$103.24	7	82	6.45	286
6	No	\$26,495.64	12	23	1.39	462
7	Yes	\$109,330.00	16	17	2.25	648
8	No	\$45,191.00	15	31	2.88	348
9	No	\$22.45	2	24	1.39	243
10	Yes	\$62,405.00	37	51	3.92	261
11	No	\$134,473.00	5	3	0.49	742
12	Yes	\$50.87	25	130	7.69	343
13	No	\$1,647.12	0	64	3.47	465
14	Yes	\$367.42	32	118	8.45	321
15	Yes	\$6,717.00	0	34	3.98	371
16	Yes	\$80.92	58	132	9.63	267
17	No	\$35.64	36	127	8.79	255
18	Yes	\$34,625.23	25	54	2.54	416
19	Yes	\$2.84	39	132	8.64	96
20	No	\$9.45	34	147	6.78	164

Table 4. Example Dataset

4.2 Example Problems

In this section we solve the 20-item example problem provided in the example dataset shown in Table 4. We start by first making the introductory level computations such as scaling by lead time demand, calculating the holding costs, shortage costs, etc. First, we must define the parameter's that will be used throughout the example problems. Note that the same parameters will be consistent throughout all example problems. The values are shown in Table 5.

Holding Cost	20%
Shortage Cost	\$10,000.00 + Unit Cost
Salvage Value	70%
Set Up Cost	\$1,000.00

Table 5. (s, S) Example Problem Parameters

We assume that the demand parameters (mean and standard deviation) are given as yearly demand and that storage cost is given as a percentage on a yearly basis. Thus, we must first scale demand to lead time, this is done by simply employing equations 4 and 5. We will demonstrate this on the first part below:

$$\bar{x}_{leadtime} = 129_{units/year} * \left(\frac{87_{days/leadtime}}{365_{days/year}} \right) \quad (43)$$

$$\bar{x}_{leadtime} = 30.75_{units/leadtime}$$

$$s_{leadtime} = 6.25_{\sqrt{units/year}} * \sqrt{\frac{87_{days/leadtime}}{365_{days/year}}} \quad (44)$$

$$s_{leadtime} = 3.05_{\sqrt{units/leadtime}}$$

From there we calculate the respective costs by using the parameters in Table 4. First, we calculate storage cost, because the storage cost is given on a yearly basis, we must scale it to lead time. Again, the first part is used to show this process below:

$$Storage\ Cost_{dollars/leadtime} = 20\%/year * \left(\frac{87 \frac{days}{leadtime}}{365 \frac{days}{year}} \right) * \$3.23 \quad (45)$$

$$Storage\ Cost_{dollars/leadtime} = \$0.15/leadtime$$

Now we calculate the salvage value by simply multiply the percentage shown in Table 5 by the unit cost. Following this we have everything that we need to calculate the holding cost per equation 3. Lastly, calculate the penalty cost by adding \$10,000 to the unit cost as shown in Table 4. Due to the simplicity of these calculations, an example is not provided, however, results are shown below in Table 6.

Item #	BGS?	Unit Cost	Inventory Level	Mean	Std	PLT (days)	Sclaed_Mean	Scaled_Std	Alpha	Beta	Holding Cost	Penalty Cost
1	No	\$3.23	4	129	6.25	87	30.75	3.05	101.40	3.30	\$0.15	\$10,003.23
2	Yes	\$96,444.00	0	7	0.52	648	12.43	0.69	319.99	25.75	-\$33,266.57	\$106,444.00
3	No	\$500.49	0	97	4.57	212	56.34	3.48	261.67	4.64	\$58.14	\$10,500.49
4	Yes	\$28,209.00	45	43	3.67	342	40.29	3.55	128.63	3.19	-\$14,460.01	\$38,209.00
5	Yes	\$103.24	7	82	6.45	286	64.25	5.71	126.64	1.97	-\$56.09	\$10,103.24
6	No	\$26,495.64	12	23	1.39	462	29.11	1.56	348.74	11.98	\$6,707.39	\$36,495.64
7	Yes	\$109,330.00	16	17	2.25	648	30.18	3.00	100.94	3.34	-\$37,711.36	\$119,330.00
8	No	\$45,191.00	15	31	2.88	348	29.56	2.82	110.12	3.73	\$8,617.24	\$55,191.00
9	No	\$22.45	2	24	1.39	243	15.98	1.13	199.73	12.50	\$2.99	\$10,022.45
10	Yes	\$62,405.00	37	51	3.92	261	36.47	3.32	120.93	3.32	-\$34,758.73	\$72,405.00
11	No	\$134,473.00	5	3	0.49	742	6.10	0.69	77.06	12.64	\$54,673.41	\$144,473.00
12	Yes	\$50.87	25	130	7.69	343	122.16	7.45	268.56	2.20	-\$26.05	\$10,050.87
13	No	\$1,647.12	0	64	3.47	465	81.53	3.91	434.41	5.33	\$419.68	\$11,647.12
14	Yes	\$367.42	32	118	8.45	321	103.78	7.92	171.50	1.65	-\$192.57	\$11,367.42
15	Yes	\$6,717.00	0	34	3.98	371	34.56	4.02	74.01	2.14	-\$3,336.42	\$16,717.00
16	Yes	\$80.92	58	132	9.63	267	96.56	8.24	137.44	1.42	-\$44.81	\$10,080.92
17	No	\$35.64	36	127	8.79	255	88.73	7.35	145.84	1.64	\$4.98	\$10,035.64
18	Yes	\$34,625.23	25	54	2.54	416	61.55	2.71	514.50	8.36	-\$16,345.01	\$44,625.23
19	Yes	\$2.84	39	132	8.64	96	34.72	4.43	61.39	1.77	-\$1.84	\$10,002.84
20	No	\$9.45	34	147	6.78	164	66.05	4.54	211.22	3.20	\$0.85	\$10,009.45

Table 6. (s, S) Example Problem Cost Calculations

We now have everything we need to calculate examples for the three different mathematical models. First, we will start with the newsvendor model as the calculation from this step is an input into the other two models. Following the newsvendor, we will look at the equation-based method and lastly, the algorithm-based method.

4.2.1 Newsvendor Example

As mentioned earlier, due to the simplicity of the newsvendor model all calculations can be made in Excel with ease, using the statistical functions built in Excel. The first step is to calculate the critical ratio or optimal service level that is shown in the right-hand side of equations 11. An example for part 1 is shown below:

$$\text{Critical ratio} = \frac{10,003.23 - 3.23}{10,003.23 + 0.15}$$

$$\text{Critical Ratio} = 0.99966 = \%99.966$$

As we see in this case, critical ratio or optimal service level is over 99%. This is since the unit cost and storage cost of the particular item is very small with respect to the penalty cost. From here we calculate the optimal number of parts to have on hand. To do this we need to specify which distribution we are assuming demand follows. For the reasons discussed in section 2.3.1 we hypothesize that the Gamma distribution is a better fit for our application. We will show results using both the Normal and Gamma distributions here for the newsvendor model and later try to attempt to answer this question in the cost benefit analysis. In the case that we assume

demand follows the Gamma distribution, we must first calculate our shape and rate parameters by using the method of moments, defined by equations 1 and 2. Note that these calculations are made with the scaled mean and scaled standard deviation parameters. After specifying the parameters, we call the GAMMA.INV function built into Excel with probability given by the critical ratio, the shape parameter, and $1/\beta$ parameter. This is because Excel uses the alternative Gamma distribution formulation that calls for shape and scale parameters. We show this in the below formula to prevent any confusion:

$$= \text{GAMMA.INV} \left(\text{Critical ratio}, \alpha, \frac{1}{\beta} \right) \quad (46)$$

For the Normal distribution, we simply call the NORM.INV function with the critical ratio, scaled mean, and scaled standard deviation. If the value calculated from the inverse distribution function is less than the inventory on hand x , the optimal inventory level y^* is brought up to x , per the constraint listed in the model. Recall that because we are approximating with continuous distributions, we must round any non-integer solutions for y^* up to the nearest integer. The only thing left to do at this point is to calculate the number of units that need to be procured λ , to ensure we are at an optimal level; this is done by applying equation 12. Results are shown in Table 7 below.

Item #	Inventory Level	Alpha	Beta	Critical Ratio	inv_Normal	inv_Gamma	y*_Norm	y*_Gamma	λ Norm	λ Gamma
1	4	101.40	3.30	0.99966	41.13	42.20	42	43	38	39
2	0	319.99	25.75	0.13665	11.67	11.67	12	12	12	12
3	0	261.67	4.64	0.94709	61.97	62.09	62	63	62	63
4	45	128.63	3.19	0.42107	39.58	39.48	45	45	0	0
5	7	126.64	1.97	0.99531	79.08	80.05	80	81	73	74
6	12	348.74	11.98	0.23147	27.97	27.96	28	28	16	16
7	16	100.94	3.34	0.12252	26.69	26.73	27	27	11	11
8	15	110.12	3.73	0.15672	26.72	26.72	27	27	12	12
9	2	199.73	12.50	0.99746	19.15	19.33	20	20	18	18
10	37	120.93	3.32	0.26563	34.39	34.33	37	37	0	0
11	5	77.06	12.64	0.05021	4.96	5.00	5	6	0	1
12	25	268.56	2.20	0.99752	143.11	144.16	144	145	119	120
13	0	434.41	5.33	0.82872	85.25	85.24	86	86	86	86
14	32	171.50	1.65	0.98282	120.54	121.23	121	122	89	90
15	0	74.01	2.14	0.74735	37.24	37.14	38	38	38	38
16	58	137.44	1.42	0.99640	118.69	120.15	119	121	61	63
17	36	145.84	1.64	0.99595	108.18	109.40	109	110	73	74
18	25	514.50	8.36	0.35360	60.53	60.49	61	61	36	36
19	39	61.39	1.77	0.99990	51.20	53.66	52	54	13	15
20	34	211.22	3.20	0.99897	80.06	80.94	81	81	47	47

Table 7. (s, S) Newsvendor Example Problem

We see that in most cases, the Gamma distributions produces the same results as the Normal distribution. The only instances when the Gamma distribution produces a different result are those when the critical ratio is high. This is not surprising as the Gamma distribution typically has a longer right tail than that of the Normal distribution. The Gamma distribution also performs very similar to the Normal distribution when the shape parameter is large. In the example problems presented here, the alpha parameter is always large, with the smallest alpha being roughly 61.

4.2.2 A Note on (s, S) Policies

Using the Gamma distribution is more computationally demanding especially when alpha values are high. Referencing the probability distribution function for the Gamma distribution in equation 28, we see that it also includes the Gamma function which is shown in equation 29. Looking at the

definition of the Gamma function we see that as alpha increases, the Gamma function increases very rapidly. This idea is presented in Table 8 below.

x	Gamma (x)
1	1
5	24
10	362880
20	1.2165E+17
50	6.0828E+62
100	9.333E+155

Table 8. Gamma Function Example

We see that the gamma function returns very high values when being evaluated at any value greater than roughly 10. Because the smallest value of alpha for which we evaluate the Gamma function at in the example problems is roughly 61, this significantly increases the computational complexities in the (s, S) inventory policies.

Another thing to note about the (s, S) policy is that since we are ordering parts for a final product that is only going to be produced for another period, keeping the inventory levels on a (s, S) policy probably isn't the most practical approach. As the periods which we are using are lead time for the given item, and demand is assumed to be zero one lead time from which the inventory levels are to be determined. Thus, it is more appropriate to pull inventory off the (s, S) policy currently and instead procure a final quantity of the item, as is the case in the newsvendor inventory model. It is for this reason that the newsvendor model is the inventory model chosen to solve the problem presented in this thesis. However, the (s, S) policies presented show potential for cost savings if implemented within a few periods of the end of a production run and then

switching to inventory levels produced by the newsvendor model one period away from the end of production. The (s, S) policies will still be demonstrated on the example problems to give the reader a demonstration and leave opportunity for future research. However, due to the computational complexities discussed earlier in this section and because our focus is on the newsvendor inventory model, only the Normal distribution will be used for demand.

4.2.3 (s, S) Equation Example

Now we will explore the equation-based method for the (s, S) policy on the example problems. The inputs for the model have previously been calculated, recall that for both the (s, S) policies, the newsvendor solution is an input into the model and the other parameters were calculated in section 4.2. However, the solution that we arrived at with our tailored version of the newsvendor model, with initial inventory considered, set a lower bound on y^* such that $y^* \geq x$. In this application we don't want that lower bound, so we simply take the original value returned from our inverse function and round up. Recall that in this model, this represents the S or max inventory level. To solve for s , we find the smallest value which satisfies equation 24. As mentioned in section 3.5, MATLAB will be used to compute the inventory levels as these policies require significantly more computation than the newsvendor model. We will set up the equation which solves for the s inventory level for part one below, although, it will not be solved by hand:

$$\begin{aligned}
& 3.23 * s + \int_s^{\infty} 10,003.23(z - s) \frac{1}{3.05\sqrt{2\pi}} * e^{-\frac{1}{2}\left(\frac{z-30.75}{3.05}\right)^2} dz \\
& + \int_0^s 0.15(s - z) \frac{1}{3.05\sqrt{2\pi}} * e^{-\frac{1}{2}\left(\frac{z-30.75}{3.05}\right)^2} dz \\
& = 1,000 + 3.23 * 38 \\
& + \int_{38}^{\infty} 10,003.23(z - 38) \frac{1}{3.05\sqrt{2\pi}} \\
& * e^{-\frac{1}{2}\left(\frac{z-30.75}{3.05}\right)^2} dz \int_0^{38} 0.15(38 - z) \frac{1}{\sigma\sqrt{2\pi}} * e^{-\frac{1}{2}\left(\frac{z-30.75}{3.05}\right)^2} dz
\end{aligned} \tag{47}$$

Solutions are obtained via MATLAB and shown in Table 9 below:

Item #	Scaled_Mean	Scaled_Std	S*	s*	s*_Rounded
1	30.75	3.05	42	35.15	35
2	12.43	0.69	12	11.12	11
3	56.34	3.48	57	59.86	59
4	40.29	3.55	41	38.60	38
5	64.25	5.71	68	73.33	73
6	29.11	1.56	28	27.46	27
7	30.18	3.00	27	25.98	25
8	29.56	2.82	27	26.02	26
9	15.98	1.13	20	17.03	17
10	36.47	3.32	35	33.40	33
11	6.10	0.69	5	4.66	4
12	122.16	7.45	137	135.09	135
13	81.53	3.91	86	83.60	83
14	103.78	7.92	117	115.67	115
15	34.56	4.02	38	35.73	35
16	96.56	8.24	109	110.88	110
17	88.73	7.35	101	101.14	101
18	61.55	2.71	61	59.64	59
19	34.72	4.43	48	41.84	41
20	66.05	4.54	79	73.27	73

Table 9. (s, S) Equation Example

Again, since we are approximating by using a continuous distribution for demand, we must round the calculated answer. In this case we round down for all the parts, to ensure a plausible policy. Note that for item 2, if the calculated answer is rounded up, both s and $S = 12$.

4.2.4 (s, S) Algorithm Example

Here the adjusted algorithm with ordering cost and new break if rule formulated in section 3.4.2.1 is demonstrated on the example dataset. Again, we use the newsvendor solution without the bound to current inventory on hand. Note that to maintain a reasonable computation time, φ is set to be 10, however, in doing this we cannot guarantee that we arrive at an optimal solution, although, the odds of arriving at the optimal solution are very likely as this is a large number of standard deviation away from the previously calculated newsvendor problem. Thus, it is highly unlikely that a better solution exists this far away from the previous solution. The results shown from running the algorithm are shown in Table 10 below.

Item #	PLT (days)	Scaled_Mean	Scaled_Std	S	S*	s*
1	87	30.75	3.05	42	42	35
2	648	12.43	0.69	12	12	11
3	212	56.34	3.48	57	62	59
4	342	40.29	3.55	41	40	38
5	286	64.25	5.71	68	80	73
6	462	29.11	1.56	28	28	27
7	648	30.18	3.00	27	27	25
8	348	29.56	2.82	27	27	26
9	243	15.98	1.13	20	20	17
10	261	36.47	3.32	35	35	33
11	742	6.10	0.69	5	5	4
12	343	122.16	7.45	137	144	135
13	465	81.53	3.91	86	86	83
14	321	103.78	7.92	117	121	115
15	371	34.56	4.02	38	38	35
16	267	96.56	8.24	109	119	110
17	255	88.73	7.35	101	109	101
18	416	61.55	2.71	61	61	59
19	96	34.72	4.43	48	52	41
20	164	66.05	4.54	79	81	73

Table 10. (s, S) Algorithm Example

Note that there is an extra S term here denoted S^* , this is because the algorithm searches for a new maximum value that improves on the newsvendor solution. In this data set, it looks like the algorithm found better solutions for S in several different occasions. Overall, the solutions look good from the algorithm and all inventory policies are plausible.

4.3 Cost Benefit Analysis

In this section we will present the cost benefit analysis conducted in attempt to estimate the monetary benefits of conducting such analysis and the resulting adjustment of inventory levels before production of a particular final product ends. A simulation based approach is taken while conducting this analysis because historical data specifically relating to the time that a product was nearing the end of production is often scarce or not available at all. It was mentioned previously

that since aerospace companies sell most of their products based on contracts, often, there is a set quantity agreed to be purchased before the items are manufactured. Due to this, there is little demand uncertainty in the number of final products that will be produced. However, raw materials, parts, and sub-assemblies that make up the final product are still subject to scrap, loss, and rework during the manufacturing of the products. This scrap, loss, and rework is the uncertainty for which we are attempting to account for during the final period of production. It is important to note that scrap, loss, or rework will never result in demand for the analyzed items to decrease as might be the case with final product demand uncertainty. In this application, the uncertainty is in the number of additional units from the MRP schedule that may be required in the case that a unit is lost, scrapped, or reworked.

The cost benefit analysis conducted here assumes there is no initial inventory on hand and that there will never be fewer units required during the last period of production than the expected demand from the MRP schedule. It seems reasonable that in certain cases, production may fall behind and continue after the point in which we are a lead time away from the expected production end date. In this case the demand will still be fulfilled, just at a later time, which may result in more storage costs to be incurred than as assumed by the model but this is assumed to be negligible at this level. However, it seems unlikely that an expected demand from the MRP schedule will not be fulfilled given the contractual nature for which the MRP schedule is produced. Noting that the goal of this cost study is not to produce an exact amount of expected cost saving but rather to give us an idea if the analysis is beneficial and to what degree. Thus, we are comfortable making this assumption and are confident that it will not have any significant impacts regarding the results found in this analysis.

This cost benefit analysis will be conducted by using the historical demand distributions and historical scrap, loss, or rework rates for a particular set of aerospace parts and sub-assemblies. Using the historical scrap, loss, or rework percentages along with the historical demand rates, a random scrap, loss, or rework quantity is produced using a random number generator that is representative of the quantity we could expect in any given lead-time adjusted period. For each part or sub-assembly for which there was a nonzero historical scrap rate, a safety stock quantity is assigned based on calculations from the newsvendor model. For the parts which had a zero historical scrap, loss, and rework rate safety stock was set to zero. Roughly 90% of the parts in the dataset used had 0% historical scrap, loss, and rework rate and thus, safety stock values were only calculated for approximately 10% of the parts.

The parameters for the historical demand distribution along with the respective costs relating to each part/sub-assembly are inputted into the newsvendor model to compute the policies as in section 4.2.1. Note that in some cases y^* will be calculated to be lower than expected demand; in this case we simply set safety stock to zero, hence, the assumption regarding demand. In cases where y^* is calculated to be greater than the expected demand, we subtract the expected demand and take the net value as our safety stock value. This process is presented below:

- If $y^* \leq \bar{x}$ then *safety stock* = 0
- If $y^* > \bar{x}$ then *safety stock* = $y^* - \bar{x}$

Two inventory policies are generated by assuming that demand follows either a Normal or Gamma distribution. Based on the random number generated for scrap, loss, and rework and the amount of safety stock, we will either have excess inventory, zero inventory, or shortages at the end of the period. Using the holding, penalty, and unit ordering costs we then compute the total cost of the given inventory policy throughout the period. Note that the unit ordering costs are also included for the expected demand along with any additional safety stock units. This cost is compared to the cost of the baseline policy, which is simply zero safety stock or additional inventory and is representative of drawing down safety stock to zero for products one period away from ending production. Because the numbers generated are randomly generated, this process is completed for 30 iterations, recording the costs of the Gamma distributed demand policy, the Normally distributed demand policy, and the baseline policy throughout each iteration. Key results on cost savings from the analysis are summarized in Table 11 below.

	% Difference in Total Cost From Baseline
Normal	-0.7217%
Gamma	-0.7205%

Table 11. Cost Benefit Analysis Overview

We see that the percent difference of the newsvendor model from the baseline policy in regards to total inventory cost throughout the period is roughly -0.72% in either case where we assume demand follows a Normal or Gamma distribution. This implies that one could expect to save roughly 0.72% of total inventory costs throughout the last lead-time period of production by setting inventory levels to the optimally calculated values given by the newsvendor model, as opposed to drawing safety stock down to zero. While a -0.72% difference between the newsvendor and

baseline models may not seem like a large value, when analyzing a large subset of parts at a time this can add up to a significant amount. As an example, 0.72% of \$300 million is \$2.160 million, note that the \$300 million is a made up number. Nevertheless, we can see how if we are analyzing a large subset of parts a significant amount of cost savings can be incurred by applying a relatively simple mathematical model to the inventory levels rather than just drawing safety stock down.

Because the vast majority of items in this dataset (~90%) have a zero historical scrap, loss, or rework rate associated with them, the costs associated with procuring the units of inventory equal to the expected demand account for a very large proportion of the total costs. We will define the additional cost as the cost of the additional units procured based on the safety stock values plus the total holding costs incurred plus the total shortage costs incurred. This definition may be easier understood by referencing the formula below:

$$\begin{aligned} \textit{additonal cost} & & (48) \\ & = (\textit{safety stock}) * c + \textit{holding costs} + \textit{shortage costs} \end{aligned}$$

We see that this additional cost value accounts for everything except for the cost associated with procuring the number of units that are equal to expected demand. This is done to better understand the factors driving the cost savings seen in the newsvendor model. Table 12 below shows the proportions of total cost that is accounted for by additional cost.

	Percent of Total Cost That's Additional Additional Cost	
	Baseline	News vendor
Normal	0.9370%	0.2237%
Gamma	0.9413%	0.2268%

Table 12. Proportion of Total Cost Accounted for by Additional Costs

In the baseline policies, additional costs make up roughly 0.94% of the total inventory costs and in the news vendor policies, additional costs only make up about 0.22% of the total inventory costs. Using the additional cost for the baseline policy, we will look at the relative proportions of cost that are due to storage, salvages, and shortages. First we start with the policy generated by assuming demand follows a Normal distribution in Figure 18 and the policy by assuming the Gamma distribution in Figure 19.

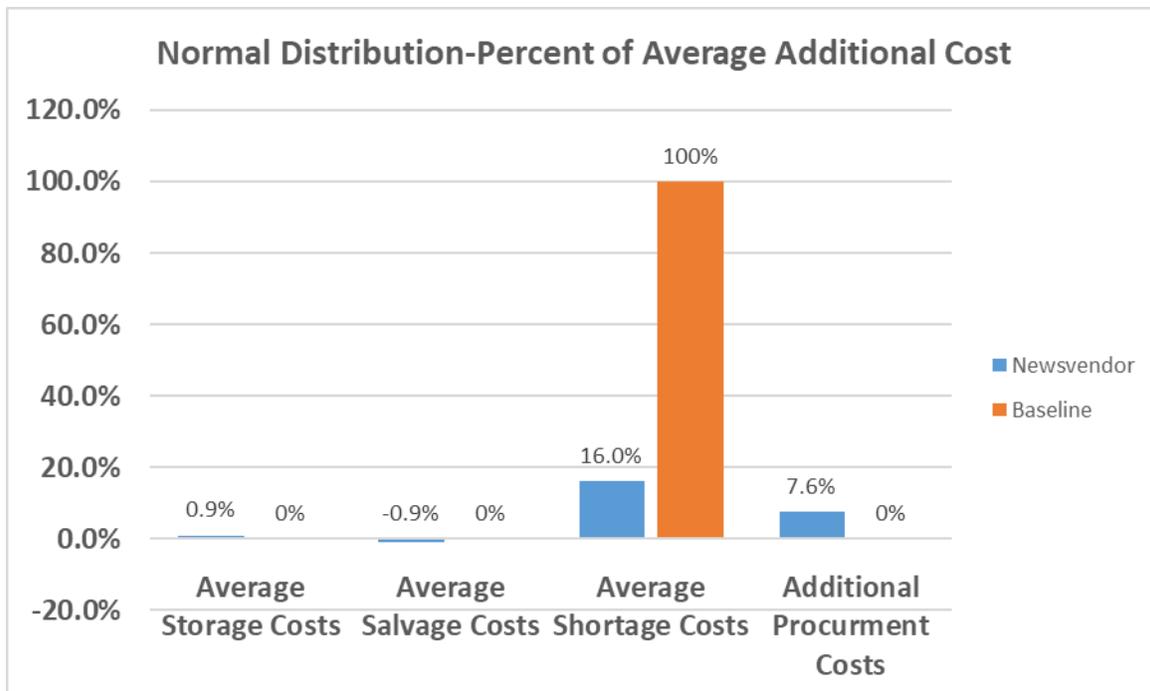


Figure 18. Additional Cost Proportions-Normal Distribution

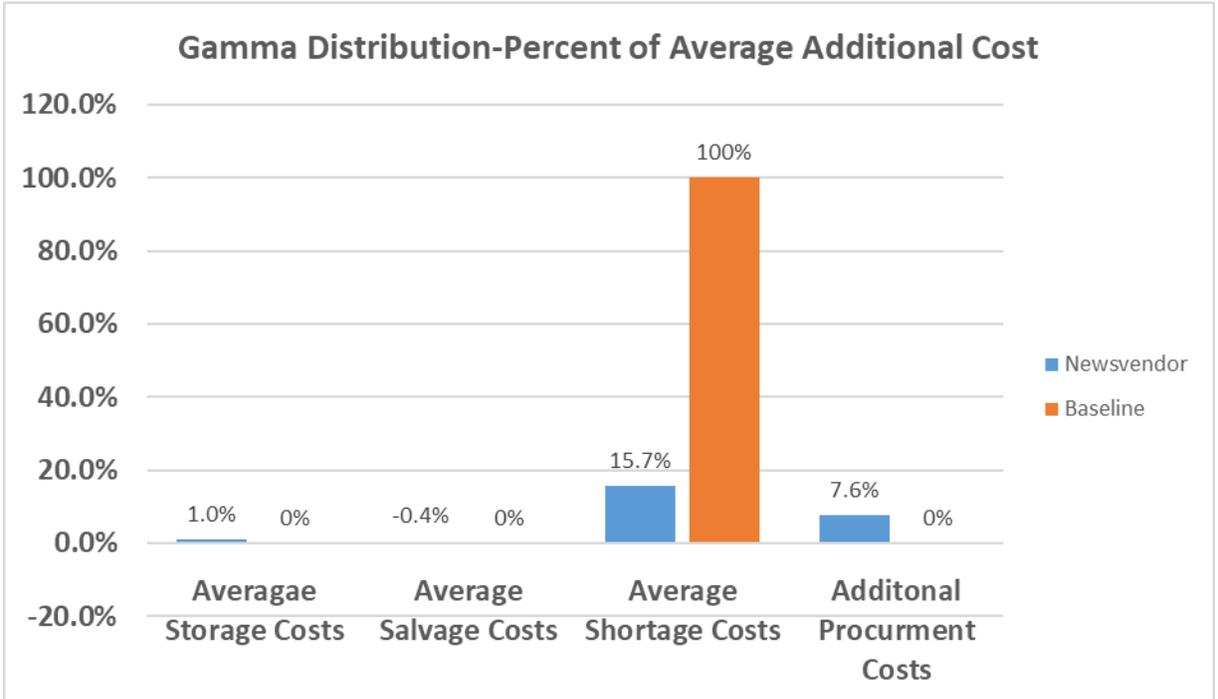


Figure 19. Additional Cost Proportions-Gamma Distribution

Recall that salvage costs are negative because salvage value represents a positive return on the initial investment if there is excess inventory at the end of the period. For the baseline case, 100% of the additional costs come from the average shortage costs, this is because no additionally inventory is procured, thus, there is no excess inventory to incur storage costs on. In the case of the newsvendor model, we see that the cost savings come from trading expected number of shortages for additional procurement, thus, increasing expected number of excess items and storage costs.

A paired sample Student's t-test is performed to test whether or not the total inventory costs for the policy produced by the newsvendor model is significantly lower than the total inventory costs

for the baseline models inventory policy. This test is completed for assuming demand follows either the Normal or Gamma distribution using a significance level of 95% in both cases. The formal hypothesis structure and tests can be seen below:

Normal Distribution:

$$H_0: \mu_{\text{Total cost of newsvendor model assuming Normal Dist.}} = \mu_{\text{Baseline model}}$$

$$H_1: \mu_{\text{Total cost of newsvendor model assuming Normal Dist.}} < \mu_{\text{Baseline model}}$$

<i>t-Test: Paired Two Sample for Means</i>	<i>Newsvendor Total Cost</i>	<i>Baseline Total Cost</i>
Mean	x	x
Variance	x	x
Observations	30	30
Pearson Correlation	0.033250387	
Hypothesized Mean Difference	0	
df	29	
t Stat	-113.1030191	
P(T<=t) one-tail	3.20765E-40	
t Critical one-tail	1.699127027	

Figure 20. T-test on Newsvendor and Baseline for Normal Distribution

Gamma Distribution:

$$H_0: \mu_{\text{Total cost of newsvendor model assuming Gamma Dist.}} = \mu_{\text{Baseline model}}$$

$$H_1: \mu_{\text{Total cost of newsvendor model assuming Gamma Dist.}} < \mu_{\text{Baseline model}}$$

<i>t-Test: Paired Two Sample for Means</i>	<i>News vendor Total Cost</i>	<i>Baseline Total Cost</i>
Mean	x	x
Variance	x	x
Observations	30	30
Pearson Correlation	-0.191742681	
Hypothesized Mean Difference	0	
df	29	
t Stat	-124.3156557	
P(T<=t) one-tail	2.08001E-41	
t Critical one-tail	1.699127027	

Figure 21. T-test on News vendor and Baseline for Gamma Distribution

Because our p-value is less than our $\alpha = 0.05$ in both the cases that demand follows a Normal or Gamma distribution, we reject the null hypothesis and conclude that the mean total cost of the inventory levels produced by the news vendor model is less than the mean total cost of the inventory levels produced by the baseline model. This is true regardless of if demand is assumed to follow a Normal or Gamma distribution. Given the results from this cost benefit analysis, we are confident that inventory policies generated from the news vendor inventory model will result in cost savings from policies that draw safety stock down to zero for the last period in production.

Now on the note of which distribution is better to use to fit demand, Gamma or Normal? We didn't see much of a difference between the two while looking at the table of results and bar charts. However, to see if a statically significant difference exists between the two a paired sample Student's t-test is conducted on the Normal distribution model and Gamma distribution model. The formal hypothesis structure and test can be seen below:

$$H_0: \mu_{\text{Total cost assuming Gamma Dist.}} = \mu_{\text{Total cost assuming Normal Dist.}}$$

$$H_1: \mu_{\text{Total cost assuming Gamma Dist.}} < \mu_{\text{Total cost assuming Normal Dist.}}$$

<i>t-Test: Paired Two Sample for Means</i>	<i>Gamma Distributed Total Cost</i>	<i>Normally Distributed Total Cost</i>
Mean	x	x
Variance	x	x
Observations	30	30
Pearson Correlation	0.229905657	
Hypothesized Mean Difference	0	
df	29	
t Stat	-1.176192879	
P(T<=t) one-tail	0.124541559	
t Critical one-tail	1.699127027	

Figure 22. T-test on Gamma and Normal Newsvendor Models

Because our p-value of roughly 0.12 is greater than our $\alpha = 0.05$, we fail to reject the null hypothesis and conclude that mean total inventory cost of the policy generated by the newsvendor model assuming demand follows a Gamma distribution is not less than the total inventory cost of the policy generated by the newsvendor model assuming demand follows a Normal distribution.

Chapter 5

Conclusion & Future Work

5.1 Conclusion

Throughout this thesis, we examined the problem of properly setting inventory levels for items nearing the end of a production run. We explained how the steady state solutions previously calculated are no longer optimal as the inventory system is no longer in a steady state and why the last period of production is such a critical time to have accurate inventory levels to fit demand. We noted how the aerospace industry makes for an interesting case of this problem due to the unique structure of their supply chains, making a shortage extremely undesirable, which isn't always the case in other industries. We discussed how the aerospace industry also typically involves little final product uncertainty, causing the variation in demand for raw materials, parts, and sub-assemblies to stem from scrap, loss, and rework throughout the manufacturing processes. The typical supply chain structure for aerospace companies was examined. Three mathematical models were tailored to fit the specific problem of determining inventory levels at the end of a production run and demonstrated. The newsvendor inventory model was the most practical as compared to the two (s, S) inventory models explored. A cost benefit analysis was then conducted on the newsvendor model using a simulated based approach. Good results were found in the cost benefit analysis and our newsvendor model was shown to perform better than the baseline model which draws safety stock down to zero. In fact, this difference was found to be statistically significant for any practical significance level for both newsvendor models that assumes demand follows a Normal or Gamma distribution. However, mean total inventory costs for the newsvendor model that assumes demand follows a Gamma distribution was not found to be significantly less than the

mean total inventory costs for the newsvendor model assuming demand follows a Normal distribution at a 95% significance level.

5.2 Future Work

Throughout completing this thesis, several different ideas for future work have been generated. The first being to further investigate the differences between using the Gamma and Normal distributions to fit demand as it is suspected that the Gamma distribution would outperform the Normal distribution with high dollar parts, which typically have low demand. Another thought for future work is to look into how penalty costs are defined and determine if the one size fits all system is appropriate. Perhaps, in future work dropping the penalty cost all together and instead modeling to service level.

Another area for future work is in using the (s, S) policies that were explored here. As discussed previously, it may be beneficial to employ these policies some number of periods out from the end of production, switching to the newsvendor model the period before production ends. The (s, S) algorithm-based method holds promise for future work. As the additional break if statement added to the formulation, reducing computation time, affects the guarantee of optimality in the solutions generated.

BIBLIOGRAPHY

Arrow, K. J., Harris, T., & Marschak, J. (1951). Optimal inventory policy. *Econometrica: Journal of the Econometric Society*, 250-272.

Boeing. <https://www.boeing.com/services/government/supply-chain-logistics.page>.

Cook, K. (2004). *TacAir team now gets components 'just in time'*. Boeing.

https://www.boeing.com/news/frontiers/archive/2004/september/i_ids1.html.

Dolgui, A., & Prodhon, C. (2007). Supply planning under uncertainties in MRP environments: A state of the art. *Annual reviews in control*, 31(2), 269-279.

Edgeworth, F. Y. (1888). The mathematical theory of banking. *Journal of the Royal Statistical Society*, 51(1), 113-127.

Forrester, J. W. (1997). Industrial dynamics. *Journal of the Operational Research Society*, 48(10), 1037-1041.

Hur, M., Keskin, B. B., & Schmidt, C. P. (2018). End-of-life inventory control of aircraft spare parts under performance based logistics. *International Journal of Production Economics*, 204, 186-203.

Hopp, W. J., & Spearman, M. L. (2004). To pull or not to pull: what is the question?. *Manufacturing & service operations management*, 6(2), 133-148.

- Keaton, M. (1995). Using the gamma distribution to model demand when lead time. *Journal of Business Logistics*, 16(1), 107.
- Lee, H. L., Padmanabhan, V., & Whang, S. (1997). The bullwhip effect in supply chains. *Sloan management review*, 38, 93-102.
- Liberopoulos, G., Tsikis, I., & Delikouras, S. (2010). Backorder penalty cost coefficient “b”: What could it be?. *International Journal of Production Economics*, 123(1), 166-178.
- McGarvey, Ronald (2019) *Stochastic Periodic-Review Models Single-Period No Setup Costs*. Personal collection of (Dr. Ronald McGarvey), University of Missouri, Columbia, Missouri
- McGarvey, Ronald (2019) *Single-Period Model with Initial Inventory and Setup Costs*. Personal collection of (Dr. Ronald McGarvey), University of Missouri, Columbia, Missouri
- Mocenco, D. (2015). Supply chain features of the aerospace industry. particular case airbus and boeing. *Scientific Bulletin-Economic Sciences*, 14(2), 17-25.
- Nahmias, S., & Olsen, T. (2015). *Production and operations analysis: strategy, quality, analytics, application*. Waveland Press, Inc.
- Parsons, J. C. W. (2004). *Using a newsvendor model for demand planning of NFL replica jerseys* (Doctoral dissertation, Massachusetts Institute of Technology).
- Scarf, H. (rep.). *The Optimality of S, s Policies in the Dynamic Inventory Problem*. Retrieved from <http://dido.wss.yale.edu/~hes/pub/ss-policies.pdf>
- Schwartz, B. L. (1966). A new approach to stockout penalties. *Management Science*, 12(12), B-538.

- Silver, E. A., Pyke, D. F., & Peterson, R. (1998). *Inventory management and production planning and scheduling* (Vol. 3, p. 30). New York: Wiley.
- Spearman, M. L., & Zazanis, M. A. (1992). Push and pull production systems: issues and comparisons. *Operations research*, 40(3), 521-532.
- Tang, C. S., Zimmerman, J. D., & Nelson, J. I. (2009, January). Managing new product development and supply chain risks: The Boeing 787 case. In *Supply Chain Forum: An International Journal* (Vol. 10, No. 2, pp. 74-86). Taylor & Francis.
- Tyworth, J. E., & Ganeshan, R. (2000). A note on solutions to the inventory model for gamma lead-time demand. *International Journal of Physical Distribution & Logistics Management*, 30(6), 534-539. doi:10.1108/09600030010340897
- Tyworth, J. E., Guo, Y., & Ganeshan, R. (1996). Inventory control under gamma demand and random lead time. *Journal of Business Logistics*, 17(1), 291.
- Veinott Jr, A. F., & Wagner, H. M. (1965). Computing optimal (s, S) inventory policies. *Management Science*, 11(5), 525-552.
- Whybark, D. C., & Williams, J. G. (1976). Material requirements planning under uncertainty. *Decision sciences*, 7(4), 595-606.
- Zheng, Y. S., & Federgruen, A. (1991). Finding optimal (s, S) policies is about as simple as evaluating a single policy. *Operations research*, 39(4), 654-665.