
The Impact of Installed Base and Machine Failure Prediction on Spare Parts
Forecasting and Inventory Planning

by:


Michael Patrick Brocks

Bachelor of Science, Supply Chain Management, Syracuse University, 2011

and

Renzo Eliseo Trujillo Castaneda

Bachelor of Science, Telecommunications Engineering,
Pontificia Universidad Catolica del Peru, 2010



SUBMITTED TO THE PROGRAM IN SUPPLY CHAIN MANAGEMENT
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

ARCHIVES

MASTER OF ENGINEERING IN LOGISTICS

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

JUNE 2016

© 2016 Michael Patrick Brocks and Renzo Eliseo Trujillo Castaneda. All rights reserved.

The authors hereby grant to MIT permission to reproduce and to distribute publicly paper
and electronic copies of this thesis document in whole or in part in any medium now
known or hereafter created.

Signature of Author..... **Signature redacted**
Master of Engineering in Logistics Program
May 11th, 2016

Signature of Author..... **Signature redacted**
Master of Engineering in Logistics Program
May 11th, 2016

Certified by..... **Signature redacted**
Dr. Daniel W. Steeneck
Postdoctoral Research Associate, Center for Transportation & Logistics
Thesis Supervisor

Accepted by..... **Signature redacted**
Dr. Yossi Sheffi
Director, Center for Transportation and Logistics
Elisha Gray II Professor of Engineering Systems
Professor, Civil and Environmental Engineering

The Impact of Installed Base and Machine Failure Prediction on Spare Parts Forecasting
and Inventory Planning

by

Michael Patrick Brocks

and

Renzo Eliseo Trujillo Castaneda

Submitted to the Program in Supply Chain Management

on May 11, 2016 in Partial Fulfillment of the

Requirements for the Degree of Master of Engineering in Logistics

Abstract

Recent advances in technological capability and economics have opened up a new world of capability known as the Internet of Things (IoT). The Internet of Things is the concept that all machines can be connected to the internet, and be remotely monitored through an infrastructure of interconnected software and hardware. Many companies are just beginning to explore the economic value that the Internet of Things can unlock, with much of the initial focus on remote diagnostics and predictive maintenance, particularly in application to industrial machines.

This research tests various scenarios of predictive failure accuracy, creating spare parts forecasts based off of varying predictive forecast parameters. We compare these scenarios and their respective outputs to a regular time-series forecasting scenario, inserting each type of forecast into a periodic review (R, S) inventory system. We measure the output of each forecast put into the system in terms of spare parts inventory levels and in-stock service performance. We find that as long as the true positive rate (TPR) and false positive rate (FPR) have different values, our model is able to hold a lower average inventory while providing a higher level of service. Additionally, as the difference between the two values increases, the average amount of inventory held decreases, while the level of service provided increases. A more detailed summary of the results found and the implications on service supply chain were developed, and further areas of research are discussed.

Thesis Supervisor: Dr. Daniel W. Steeneck

Title: Postdoctoral Research Associate, MIT Center for Transportation & Logistics

Acknowledgement

To my family, especially my parents; without you I would not be where I am today. Also, a heartfelt thanks to Steph – your support and encouragement made a challenging time much easier. - Mike

I would like to thank my parents for always supporting me unconditionally. To my brother Juan Trujillo, for always being my role model and the reason for being here at MIT. Also, a special thanks to Eliana Garate for helping me in the hardest moments and for encouraging me to follow my dreams. - Renzo T.

Both the authors also would like to acknowledge Dr. Daniel Steeneck for his guidance and time during this project. Additionally, they also wish to thank their classmates for making their time in Cambridge a memorable one.

Table of Contents

Abstract	2
Acknowledgement	3
1. Introduction	7
2. Literature Review	10
2.1. Service Supply Chain Criticality	10
2.2. Service Supply Chain Challenges	11
2.3. Service Supply Chain Structure	13
2.4. Warranties and Service Contracts	15
2.5. Spare Parts Forecasting	17
2.6. Spare Parts Inventory Control	19
2.6.1. Repairable Parts	20
2.6.2. Last Time Buy	21
2.6.3. ABC Classification	22
2.7. Machine Data and the Internet of Things	22
3. Methodology	24
3.1. Key Supply Chain Metrics	25
3.2. Part Dispatch Simulation	26
3.3. Installed Base and Demand Creation	27
3.4. Time-series Forecasting Methods	32
3.5. Machine Failure Prediction	34
3.6. General Model Design	40
3.7. Time-series Model Design	41
3.8. Predictive Forecast Model Design	43
3.9. Simulation Approach	44
4. Results & Data Analysis	46
4.1. Summary Statistics	46
4.3. Comparison of Performance Between Models	48
5. Discussion	54
5.1. Summary of Outcomes	54
6. Conclusion	56
6.1. Major findings	56
6.2. Implications on Service Supply Chain Operations	57
6.3. Extensions to our work and avenues for further exploration	59

7. Appendix.....	62
7.1 <i>Performance metrics for all combinations of TPR and FPR across fifteen simulations..</i>	62
8. Bibliography.....	66

Tables and Figures

<i>Figure 2.3.1. Adapted from: (Jalil, Zuidwijk, Fleischmann, & van Nunen, 2011).</i>	15
<i>Figure 2.4.1. Adapted from: (Murthy et al., 2004).</i>	16
<i>Figure 2.5.1. Adapted from: (Dekker, Pince, Zuidwijk, & Jalil, 2013).</i>	17
<i>Figure 2.6.1. Adapted from: (Schradly, 1967)</i>	20
<i>Table 3.1.1. Inventory metrics.</i>	26
<i>Table 3.3.1. Bass Model parameters.</i>	27
<i>Table 3.3.2. Warranty assignment parameters.</i>	29
<i>Figure 3.3.1. Installed base of machines over time.</i>	29
<i>Table 3.3.3. Failure distribution parameters.</i>	30
<i>Figure 3.3.2. Weibull failure rate function at $\beta=0.33$, $\beta=1$, $\beta=2$ & $\beta=3$.</i>	31
<i>Figure 3.3.3. Simulated demand curve based off of expected failures.</i>	32
<i>Table 3.4.1. Time-series forecasting parameters</i>	33
<i>Table 3.5.1. Binary classification matrix.</i>	34
<i>Table 3.5.2. Binary classification parameters.</i>	35
<i>Table 3.5.3. Predictive forecast parameters.</i>	38
<i>Figure 3.8.1. Periods of demand covered by R, S periodic review policy</i>	44
<i>Table 4.1.1. Standard error parameters.</i>	46
<i>Table 4.1.2. Time-series performance metrics.</i>	47
<i>Table 4.3.1. Percent reduction in inventory parameters.</i>	48
<i>Figure 4.3.1. Percent reduction in inventory at varying TPR.</i>	49
<i>Figure 4.3.5. Plot of all combinations of TPR and FPR at 0.1% FPR intervals and 10% TPR intervals.</i>	53
<i>Table 7.1.1 Average Inventory for all combinations of TPR and FPR</i>	62
<i>Table 7.1.2 Cycle Service Level for all combinations of TPR and FPR</i>	63
<i>Table 7.1.3 Item Fill Rate for all combinations of TPR and FPR</i>	64
<i>Table 7.1.4 Average Error for all combinations of TPR and FPR</i>	65
<i>Table 7.1.5 Bias Error for all combinations of TPR and FPR</i>	65

1. Introduction

Most of the focus on the field of supply chain management is on optimally sourcing, producing and distributing a product to a customer. However, the service supply chain is concerned with servicing that product after it has been sold. Accordingly, the service supply chain has a number of challenges that are unique to it, the most prominent of which is the forecasting and management of spare parts inventory. In a typical supply chain, component needs are known, being driven by end product demand exploding down to the component level via a bill of material, or BOM. However, spare parts needs in a service supply chain are largely unpredictable, being driven by product failures that are unique to each product in service (e.g., a failure in one instance might create demand for spare part A, while another failure might create demand for spare part B). However, this causal nature of demand presents an opportunity for a new forecasting method that is the basis of this thesis. Here, we propose an approach that incorporates real-time machine data and predictive failure analytics to forecast spare parts needs ahead of when they occur. Because companies have only recently begun building in real-time monitoring systems, including integrated sensors and network connections (commonly known as the Internet of Things, or “IoT”) into their products to analyze and communicate machine status information, this is a method of forecasting that has become a possibility only within the last few years.

This research compares three different spare parts forecasting and inventory control policies against one another via an inventory planning simulation model developed in Microsoft Excel. Two of these methods are based solely on time-series forecasting methods, which are combined with a periodic review ‘R, S’ inventory ordering policy to create a total inventory forecasting and ordering system. The third method we test uses a

binary classification system that allows the user of the model to assign both a specificity and sensitivity parameter in order to generate a forecast based off of a theoretical predictive analytics scenario. We evaluate this model against the performance of the other two models we test, and find the value of various prediction levels in terms of service and cost performance. We also discuss additional opportunities and potential uses of such a model, as well as possible limitations and constraints.

The motivation for this research is to ascertain how predictive analytics and the Internet of Things can fit into a company's existing inventory planning operations, and to quantify how valuable a predictive analytics approach is in comparison to a time-series forecasting approach at various prediction capabilities. Specifically, we seek to answer the following questions:

- What is the reduction in average inventory and/or increase in CSL/IFR for every interval (eg, 10%) increase in TPR and/or FPR?
- How does this model perform under other input parameters, most notably vendor lead time?
- What are the limiting assumptions and realistic conditions that might challenge such a model?
- How could a company possibly leverage a predictive analytics approach to re-think their entire service supply chain design and operations?

Our thesis sponsor company is a global service provider of post-sales supply chain operations. The company works with clients in a variety of industries by managing service requests from initial case opening to case resolution. This includes handling of all client's customer contacts, scheduling and dispatching of both spare parts and technicians to client's customer to fulfill service requests, and managing the flow of failed parts back to client repair and triage centers. To execute all of these activities, the company works closely with a number of 3PL and warehouse vendors to make the sure the right parts and repairs take place to solve a service request as quickly and accurately as possible. As such a provider, our sponsor company is constantly exploring methods in which it may be able to lower the costs of managing customer service requests. In this thesis, we focus on the spare parts planning aspect of service supply chain operations.

2. Literature Review

The relevant streams to our research work are an assessment of service supply chains in general, as well as a review of several different areas of the service supply chain, including warranties and service contracts, spare parts forecasting, spare parts inventory control, and finally machine data and the Internet of Things. This overview helps establish a baseline for our research, as well as provides several ideas that we may be able to incorporate in our own analysis of the specific problem our thesis seeks to address.

2.1. Service Supply Chain Criticality

Service supply chains are continuing to gain importance for companies. As Cohen et al. (2006) highlights, the globalization of technology and manufacturing standards has homogenized many products, increasing competition and reducing margins. Accordingly, companies are increasingly turning to after-sales service as a way to differentiate their products. An Advanced Market Research report states that while companies' post-sales service accounts for 24% of revenues, it accounts for 45% of profits. Similarly, Murthy, Solem and Roren (2004) highlights the increasing importance that warranty and post-sales support has taken on for companies, citing the typical profit margin for after-sales service is about 30%, whereas the typical margin on the initial product sale is about 10%. Finally, Cohen and Lee (1990) argues that exemplary post-sales service can lead to significant increases in first time and repeat sales, and thus that post-sales service should be an explicit part of corporate strategy.

We also find it important to highlight that the value of the reverse logistics market has been growing over time. Douthit, Flach and Agrawal (2011) estimates that companies

in the consumer electronics industry would spend \$16.7 billion in reverse logistics over the course of that year, which represented approximately 6% of the revenue of that industry. Similarly, Blumberg (1999) shows that the total U.S. market for reverse logistics grew from \$4.7 billion in 1996 to an estimated size of \$7.7 billion in 2000.

2.2. Service Supply Chain Challenges

In order to realize the potential in service supply chains, the cost of running it must be minimized. McGuire (1980) argues that warranty servicing costs can range anywhere from 1% to 10% of sales revenue. As Jones (2014) details, Apple's warranty servicing costs were 2.7% of revenue for both the 2013 and 2014 fiscal year, amounting to \$4.6 and \$4.9 billion respectively.

Cohen, Agrawal and Agrawal (2006) highlights the challenges of managing service supply chains. They write of the need to support not only a company's current line of products, but also the products that it has sold in the past. Given that each generation of products can use different components and vendors, the service network often has to manage up to 20 times the number of stock keeping units, or SKUs, that a regular manufacturing function needs to. As discussed in the introduction, there is also inherent demand unpredictability in service supply chains, given that repairs occur sporadically and inconsistently. Finally, the challenge of returning, testing, repairing and/or disposing of failed products is also discussed. To successfully deal with these challenges, three resources that are critical to coordinate within a service supply chain are highlighted:

- material (spare parts),

-
- people (call center staff, repair depot and warehouse staff, and transportation staff),
 - infrastructure (information systems, material handling equipment, communication systems, and repair and testing equipment).

The SKU complexity problem highlighted by Cohen et. al is amplified by what is known as the “lifecycle mismatch” problem. Solomon, Sandborn and Pecht (2000) highlights this particular difficulty with spare parts management in the service supply chain, citing the difference between the increasingly short manufacturing lifecycles of products and the relatively stable or even increasing warranty periods that are provided with them and demanded by consumers. This difference creates a mismatch between when components are manufactured and/or available for procurement versus when they are demanded for replacement or repair. It’s suggested in the research that the primary way to deal with this issue is through a last time buy, which we review in Section 2.6. We also point the reader to Fleischmann, van Nunen and Gräve (2003) for an analysis of IBM’s use of repaired components in their spare parts inventory. The research highlights how IBM’s Global Asset Recovery Division exploits the relatively longer lifecycle of components (as compared to the end product) to deal with the lifecycle mismatch problem, using returned products from the field to help restock its spare parts inventories. The cost advantages of such a practice are also discussed. Berger (2012) highlights the costs of handling spare parts inventories, including the cost of capital on the capital investment, ordering costs, holding and material handling costs, and insurance, obsolescence and shrinkage costs. Cohen et al. (2006) expands on this last cost, stating that while original equipment manufacturers (OEMs) carry spares worth 10% of their annual sales on average, 23% of

parts become obsolete every year, and inventory turns hover at just one to two times annually.

Finally, one of the biggest challenges in the service supply chain, especially in the consumer electronics industry, is the phenomenon of No Fault Found (NFF). Douthit et al. (2011) found that 68% of products returned in the consumer electronics industry were found to have no fault in their performance upon inspection. As Overton (2005) highlights, in the worldwide mobile industry alone, NFF cost companies \$4.5 billion, or roughly \$55 per returned device for administrative, logistics and refurbishment costs. Qi, Ganesan and Pecht (2008) describes some of the reasons for the high prevalence of NFF, grouping reasons into the categories of people and skills, communication of the problem, testing equipment issues and methods, and finally, intermittent electronic failures.

2.3. Service Supply Chain Structure

Service supply chain activities can be classified into pre- and post-repair operations. Pre-repair operations encompass the forecasting, purchasing, transportation and warehousing of spare parts, handling and recording customer service requests, and the scheduling and dispatching of parts and field service engineers to physically perform the system repair. Post-repair operations, or reverse logistics, encompass what Guide, Harrison and Van Wassenhove (2003) summarizes into five key functions; product acquisition, reverse logistics, test and disposition, repair and/or refurbishment of the failed component, and marketing for reuse and/or resale.

Cohen and Agrawal (1999) categorizes service levels in the computer industry into two different buckets, “enterprise” and “non-enterprise.” It’s explained that enterprise

customers often have service response requirements around one to eight hours, while non-enterprise customers service requests typically require next business day (NBD) service. As Cohen and Lee (1990) points out, these service categories also affect whether companies choose to service products in-house or in the field. In-house service requires the end products to be transported back to a repair facility and is generally used for less expensive, more mobile products, such as a personal computer. Alternatively, field service entails spare parts components and a field engineer being dispatched to a customer site, which is typically used for expensive, stationary products, such as a server or industrial equipment.

Cohen and Lee (1990) emphasizes the particular attention that needs to be paid to the design and control of the service supply chain network. The research argues that the location, capacity, capability and material flow through each facility in the network is dependent on the product attributes, required service response time, and geographic dispersion of the customer set for a product. The research notes that the required service level is reflected by the number of echelons in a service supply chain network, which typically ranges from two to five. Cohen et al. (2006) refers to this decision as the “hierarchy of geography”, or deciding at what geographic distance from the customer to stock spare parts. Figure 2.3.1 shows how the different service requirements and geographic dispersion of IBM’s customers drive the location of their spare parts stocking locations in a field repair service network.

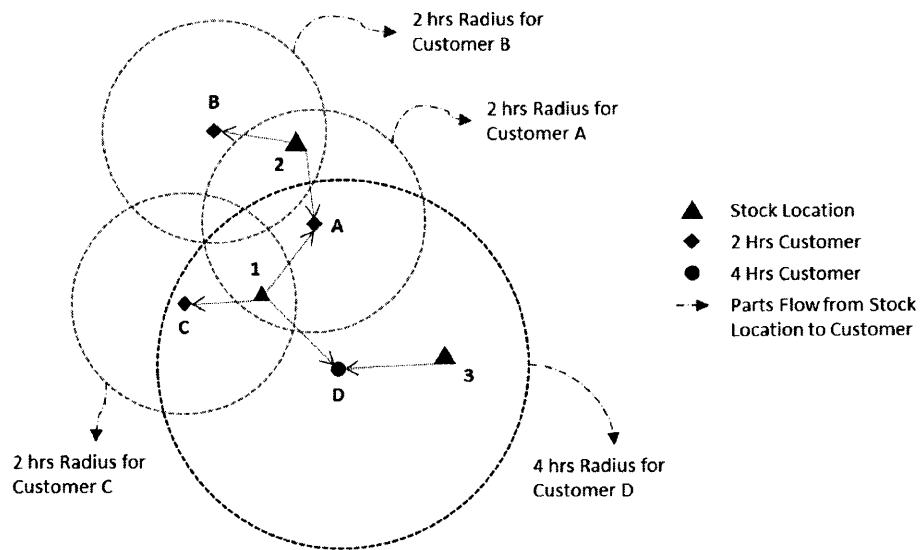


Figure 2.3.1. Adapted from: (Jalil, Zuidwijk, Fleischmann, & van Nunen, 2011).

2.4. Warranties and Service Contracts

Priest (1981) highlights the use of a warranty as both a contract for repair and an insurance policy for the consumer. Furthermore, the use of warranties as signal to the consumer of the related product's performance is also highlighted, and thus the possible use of warranties as a marketing tool. Blischke and Murthy (1992) explains both the reasons and economics behind different warranty and service contract offerings, as well as provides a framework for their categorization. Murthy et al. (2004) considers the impact that warranties have on the location of facilities that comprise a reverse supply chain, which is illustrated in Figure 2.3.1. Additionally, they also synthesize the different causes of warranty claims, shown in Figure 2.4.1, and illustrate how the number of claims/service requests drives the need for proper forecasting and positioning of spare parts in the service supply chain network. We expand on these points in the next two sections, Spare Parts Inventory Forecasting (2.5) and Spare Parts Inventory Control (2.6).

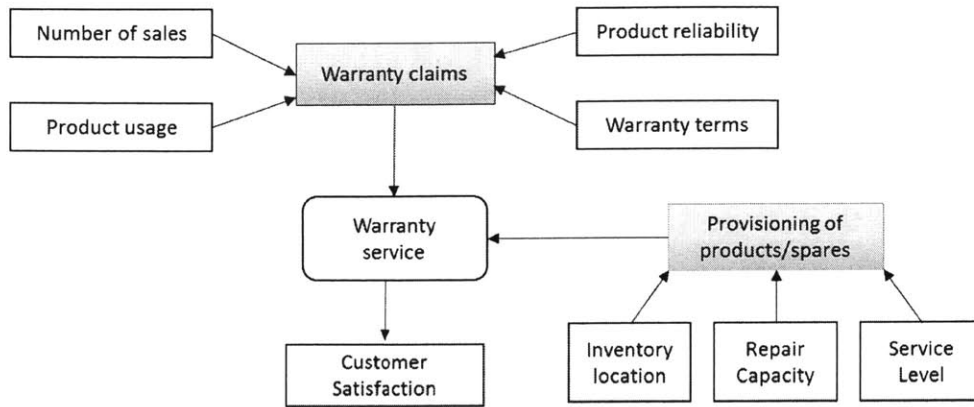


Figure 2.4.1. Adapted from: (Murthy et al, 2004).

Warranties in the consumer electronics industry are typically offered in different tiers. For example, Apple offers its MacBook product with a standard one-year hardware replacement or repair warranty and 90 days of phone support. However, the company also provides the option to upgrade to a three-year extended warranty (“AppleCare Plus”) that extends both the hardware and phone support components of the warranty (“AppleCare Protection Plan for Mac,” n.d.). Other consumer electronics companies, such as Toshiba (“Toshiba Extended Service Plans,” n.d.) offer similar warranties. Companies also have varying service period lengths for which they will offer hardware support. For example, Apple offers hardware repair and replace services for up to seven years after the date of manufacture, after which it considers products obsolete (“Apple Vintage and Obsolete Products,” n.d.).

2.5. Spare Parts Forecasting

Effective spare parts inventory control and satisfactory customer service rely heavily upon the accuracy of the forecasted demand. In the following section, we review a variety of traditional forecasting techniques for spare parts. The chosen forecasting technique is most dependent on the frequency of spare part demand, as well as the stage of the lifecycle that the product is in. Figure 2.5.1 shows how the lifecycle of the product drives the demand for spare parts.

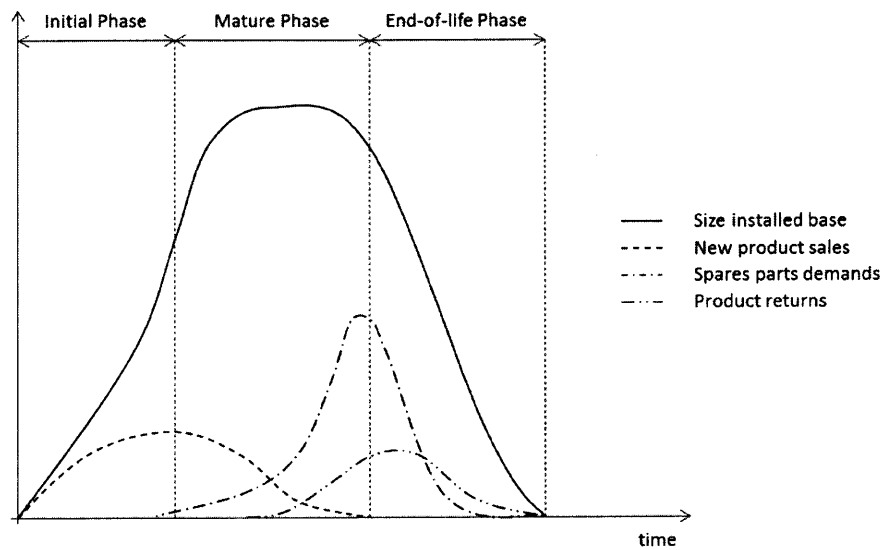


Figure 2.5.1. Adapted from: (Dekker, Pince, Zuidwijk, & Jalil, 2013).

To forecast spare parts demand during the initial phase, companies often forecast sales and expected failures for some given time period based on data from similar product launches, and multiply those two numbers to come up with an expected number of spare parts required. Expected failures are typically based on reliability engineering data, while new product sales are often based on the Bass diffusion model (Bass, 1969). Given the lack of other options for forecasting at this stage of the lifecycle, we end our review of such methods here.

Time-series forecasting is one of the most frequently used forecasting methods for the mature phase of the product lifecycle, with some of the most commonly used methods being a family of techniques known as exponential smoothing. Holt (1959) highlights the mechanics of how exponential smoothing works, and identifies six different variants of the exponential smoothing model to adjust a forecast for the presence of both seasonality and trend. We also point the reader to Winter (1960) for additional research on the exponential smoothing method, and Brown (1959) for a review of various forecasting methods, error measurements and corresponding level of service calculations. This literature is critical to our work, as it describes the mechanics of a time-series forecasting approach, which we use as a baseline to compare the performance of our predictive analytics model to.

However, spare parts demand is often slow moving, or intermittent. Observing that the use of the exponential smoothing method proposed by Holt often leads to unnecessarily high stock levels for items with this type of demand, Croston (1972) suggests an improved time-series forecasting process by modifying the exponential smoothing method. He bases his model on two separate factors; the average inter-arrival time of demand and the average size of demand, dividing these two numbers to come up with an average demand level to use in an exponential smoothing model. In 2001, Croston's method was corrected for the bias that existed in the original model for different smoothing constant values (Syntetos & Boylan, 2001). Willemain, Smart and Schwarz (2004) proposes an alternate method for forecasting intermittent demand based on the bootstrapping method. In this method, the authors account for demand autocorrelation across successive periods, and create their forecast for the non-zero periods through random sampling of the industrial dataset they used in their study.

Once a product enters the decline stage of its lifecycle, a different forecasting method can become more appropriate. Such a method is that proposed by Moore (1971), which uses decay curves to model the decline in spare part demand. Three variants of decay curves are identified, with the one that best fits the existing end product sales data on a logarithmic scale being selected. This decay curve is then transformed from a logarithmic to arithmetic scale to create a forecast for future spare parts demand during the end of the lifecycle. Ritchie and Wilcox (1977) proposes another method for forecasting end-of-lifecycle demand based on renewal theory. This method creates a demand for spare parts based on a probability of failure over some unit of time, the number of machines in service, and the probability in some period that a failed part will be replaced. This research is interesting to us due to its similarity to our research in using failure probabilities and the size of the installed base in order to create a forecast for spare parts demand.

2.6. Spare Parts Inventory Control

Spare parts inventory control presents a unique challenge in that inventory is available from two sources; procuring new spare parts from the OEM vendor, or successfully repairing failed spare parts that have come back from the field in a previous service request. Here, we review several methods on the optimal management of such a scenario.

2.6.1. Repairable Parts

Schrady (1967) proposes optimal order policies for a single-echelon inventory repair network; both the quantity of new spare parts to procure from a vendor and the quantity of failed parts to order from a repair depot are determined, given the probability of repair of the failed spare parts arriving to minimize the total cost function. A deterministic arrival rate of repairs with non-zero, deterministic lead and repair time and no backorders allowed is assumed in a simple network of a customer, a repair center and a warehouse. Two different policies are proposed; the first policy calls for a repair center inducing a batch of “not ready for issue” (NRFI) parts once the NRFI inventory builds to a certain level. The second policy uses the “ready for issue” (RFI) inventory dropping to a certain level to call for the induction of another batch of NRFI inventory into the repair process. Once there is not enough NRFI inventory to reach the required batch size for repair, a batch of new parts is procured. Figure 2.6.1 illustrates how the second policy works. Q_R denotes the repair batch size, while Q_P indicates the procurement batch size.




Figure 2.6.1. Adapted from: (Schrady, 1967)

Allen and D'Esopo (S. Allen & D'Esopo, 1968) builds on Schrady's model; however, the timing of when to order parts through a repair operation, rather than considering the optimal amount of parts to repair or procure, is the focus of the research. A Poisson arrival rate of demand, the allowance for backorders and a non-zero lead time are all assumed, with the goal of the model to minimize holding, shortage and ordering costs in a repair network.

2.6.2. Last Time Buy

As highlighted through our discussion of the lifecycle mismatch problem in Section 2.2, one of the more challenging problems in spare parts inventory planning is dealing with obsolescence. One such inventory control model that deals with the obsolescence problem is that proposed in Pince and Dekker (2011), which develops an optimal order policy for a single-item, single-echelon inventory network to incorporate the risk of obsolescence. This model proposes a shift in the base stock policy ahead of a product entering into obsolescence, allowing the remaining demand to consume the existing inventory before entering into the partial or full obsolescence period. Across all of the scenarios tested, it's found that on average, the cost of not adopting the proposed policy is roughly a 30% increase above the minimum cost found using the model. Alternatively, Fortuin (1980) considers an optimal last time buy policy where an organization must decide how many spare parts to order to cover the demand in the remaining service period (RSP). Here, a mathematical function is developed that balances the availability degree against the obsolescence risk and the shortage risk assuming normally distributed and exponentially decreasing demand for non-repairable service parts. For further research on the last time

buy problem, we refer the reader to Teunter and Fortuin (1999) and Teunter and Haneveld (2002).

2.6.3. ABC Classification

One other general method of controlling inventory that we wish to mention in this section is the ABC classification method. While it is one of the most popular methods for controlling finished goods inventory, there has been relatively little research to its application to spare parts inventories. Flowers and O'Neill (1978) considers the application of an ABC analysis to spare parts inventory at a factory in Texas, grouping parts into three different categories based on their frequency of use and their monetary value. Similarly, Molenaers, Baets, Pintelon and Waeyenbergh (2012) considers the classification of spare parts at a European industrial company, but base their classification on the criticality of the item rather than its monetary value and frequency. We find this last approach especially appealing to our thesis given its focus on qualitative factors of categorization, which we believe can be a useful approach to classifying parts based on the ability to predict their failure in advance.

2.7. Machine Data and the Internet of Things

As highlighted in the introduction section, the inspiration for this thesis comes from the relatively recent ability of machines to communicate to one another. Porter and Heppelmann (2015) highlights the potential of smart, connected devices (commonly known as the "Internet of Things", or IoT), describing how embedded sensors in products can now wirelessly communicate information such as their condition, location,

environment and usage, which software programs can analyze and turn into insights. These recent technological developments are particularly appealing in the area of service supply chain management. For example, Louit, Pascual, Banjevic and Jardine (2011) considers the ordering of spare parts based on the conditional monitoring of the system they are components of, using a reliability function to trigger the ordering of spares and maintenance actions once the risk of a failure reaches some unacceptable level. Li and Ryan (2011) utilizes a similar approach to consider the management of spare parts in a periodic review system. An exponential time to failure is assumed, creating a Poisson demand distribution. This demand distribution is updated as new degradation signals from machine sensors become available, creating an adaptive spare parts forecast. Both an individual machine inventory control policy, as well as an inventory control policy for a system of machines, is tested. Godoy, Pascual and Knights (2013), Elwany and Gebraeel (2008), as well as Tract, Goch, Schuh, Sorg and Westerkamp (2013), also provide approaches to the conditional based ordering of spares problem. However, we believe that to date no has analyzed the inventory and service-level impact that different levels of prediction accuracy can have on a service supply chain's operations.

3. Methodology

The overall effort of our project is to determine the value of additional information in the spare parts planning and dispatching processes. The two specific types of additional information we would like to quantify are the size of the installed base, and the probability of failure given some signal from machine data, which we use a binary classification model to analyze. A critical assumption of our project is that an organization who wishes to apply the methods and accompanying results discussed later in this research already have the necessary hardware infrastructure and software and analytical capabilities to be able to:

- accurately measure their installed base of some product over time, and
- predict the failure of hardware with some level of advance notice and accuracy.

We separate our approach to the problem into two different categories. The first category, which the methodology and results sections focus on, deals with the centralized procurement function and the overall forecasting, purchasing and planning of spare parts inventories. The second category deals with the related qualitative benefits from being able to better manage and dispatch spare parts inventory in a field service network, and is one of the core focuses of the discussion section.

To complete this analysis, we develop a Microsoft Excel based simulation framework in order to quantify the benefits of varying levels of information. To do this, we consider two different planning scenarios:

1. time-series forecasting, in which inventory ordering and dispatches are based solely off of historical trends, and

-
2. predictive failure analytics, in which inventory ordering and dispatches are based off of signals from each machine in the installed base population indicating a future failure.

3.1. Key Supply Chain Metrics

The two key sets of metrics we use to quantify the value of this additional information are service level metrics, which includes the Item Fill Rate (IFR) and Cycle Service Level (CSL) metric, as well as the average inventory held throughout the simulation model horizon, which we consider as comparable to a cost metric. Because cost data, including the cost per spare part, inventory holding rate and cost per stockout/cost per item short are not available, and because depending on the costs assumed for each could significantly influence the results of our analysis, we decided to analyze cost based simply on the average amount of inventory held.

Table 3.1.1 introduces the necessary notation and definitions to express these key metrics.

Table 3.1.1. Inventory metrics.

Notation	Description
a_i	time period with no stockout, or a time period where $IOH_{in_t} \geq d_t$
n	number of time periods through demand horizon
us_i	number of units short in time period t, or $\max(0, d_t - IOH_{in_t})$
IOH_{in_i}	inventory on hand going into time period t
IOH_{out_i}	inventory on hand coming out of time period t, if $IOH_{in_t} \geq d_t$, $IOH_{in_t} - d_t$, else 0
d_i	demand in time period t

Using the notation in Table 3.1.1, we define the equations for these metrics below.

$$Cycle\ Service\ Level = \frac{\sum_{i=1}^n a_i}{n} \quad (1)$$

$$Item\ Fill\ Rate = \frac{\sum_{i=1}^n us_i}{\sum_{i=1}^n d_i} \quad (2)$$

$$Average\ Inventory = \frac{\sum_{i=1}^n IOH_{in_i} + \sum_{i=1}^n IOH_{out_i}}{2 * n} \quad (3)$$

3.2. Part Dispatch Simulation

Our simulation model is based on the following set of steps.

1. Generate a realistic data set of sales information,
2. Assign warranty periods to each machine sold,
3. Generate the installed base using the sales and warranty information from 1 and 2,
4. Generate a demand for spare parts based on the installed base population and some failure rate, and

-
- Build the two types of forecasts (time-series and predictive) based on the demand generated from step 4.

More details about these steps are explained in Sections 3.3, 3.4 and 3.5.

3.3 Installed Base and Demand Creation

To conduct this analysis, it was first necessary to generate a realistic data set of sales information, off of which a data set of installed base information could be based. In order to generate sales information, we use the Bass Diffusion Model (1969).

Table 3.3.1 introduces some notation and accompanying definitions necessary to express the Bass Diffusion Model.

Table 3.3.1. Bass Model parameters.

Notation	Description	Assigned Value
A_t	Aggregate sales at time t	
m	Maximum size of installed base	1,000
p	Coefficient of innovation	.003
q	Coefficient of contagion	.08
D_t	Distinct sales at time t	

The following formula gives us the aggregate amount of sales through some time period t.

$$A_t = m \times \frac{1 - e^{-(p+q)t}}{1 + \left(\frac{q}{p}\right) \times e^{-(p+q)t}} \quad t = 1, \dots, n \quad (4)$$

To disaggregate this level of sales to get the amount of machines sold in each individual period, we used the following formula:

$$D_t = A_t - A_{t-1} \quad t = 1, \dots, n \quad (5)$$

We chose m to be 1,000 so the installed base size is nontrivial while the model is computational tractable. Additionally, p and q values were each chosen based off of realistic values found from analysis of previous industry sales data as shown in Bass (1969). After generating the sales data for the installed base, we then assign a unique machine id to each machine in the population. This data set gives us the time period at which a machine entered the installed base population, as well as the unique id for that machine which we use to identify it in the simulation model.

Next, we generate the time period at which a machine will exit the population. We define a machine exiting the population as its warranty period expiring, even though the machine may still be in service and require spare parts in the future. That scenario is outside the scope of this thesis, as we are concerned only with spare parts support for machines under warranty in each period. To do this, we create a set of warranty periods and proportion of the population each warranty period would be assigned to. The two warranty periods we select and their proportions are shown in Table 3.3.2. The periods and proportions were both assigned arbitrarily in order to reflect varying installed base sizes over time and differing consumer preferences.

Table 3.3.2. Warranty assignment parameters.

Warranty ID	Warranty Length (days)	Warranty Proportion of Population	Warranty Cumulative Proportion
A	156	0.3	0.3
B	260	0.7	1.0

Warranty types are randomly assigned to each machine in the installed base according to the warranty proportions in Table 3.3.2. By adding the date of each machine's sale and its warranty period together, we obtain the date of exit for each machine in the population. Finally, we obtain the installed base, without expired-warranty machines, by counting the number of machines that have not achieved the date of exit for each time period t .

Figure 3.3.1 shows the installed base over time based on the parameters shown in Tables 3.3.1 and 3.3.2.




Figure 3.3.1. Installed base of machines over time.

The next step in the process is to generate a demand for spare parts based on the installed base population. We use the Weibull distribution to simulate the probability of a machine in the installed base failing during some period of time.

Table 3.3.3 introduces some notation with definitions necessary to express the Weibull distribution.

Table 3.3.3. Failure distribution parameters.

Notation	Description	Assigned Value
a_{j_t}	Age of machine j at time t	$t - s_{j_t}$
s_{j_t}	Sale date of machine j	Given by Bass model
β	Shape parameter	1
η	Mean time to failure (MTTF)	180
γ	Location parameter	0

The cumulative distribution function (CDF) of the Weibull distribution is defined below.

$$F_{j_t} = 1 - e^{\left(\frac{-(a_{j_t}-\gamma)}{\eta}\right)^\beta} \quad t = 1, \dots, n, j = 1, \dots, m \quad (7)$$

The Weibull failure rate function over time for a machine in the installed base population sold at any time t is given by:

$$\lambda_t = \frac{\beta}{\eta} \times \left(\frac{a_{j_t}-\gamma}{\eta}\right)^{\beta-1} \quad t = 1, \dots, n \quad (8)$$

The advantage of using a Weibull distribution is that it allows for the control of different rates of failure over a period of time through the use of varying levels of β . A β value < 1 indicates that the failure rate is decreasing over time, a β value $= 1$ means a constant level of failure over time, while a $\beta > 1$ means an increasing level of failure over time. Note that at $\beta=1$, the Weibull distribution becomes equivalent to the exponential distribution with constant failure rate over time. Thus, by setting the value of β in our failure model to 1, we assume:

- machines of any age in our population have the same probability of failure, and
- the distribution of time to failure follows an exponential distribution.

In Figure 3.3.2, we show the Weibull failure rate function for $\gamma = 0$ at varying β values.

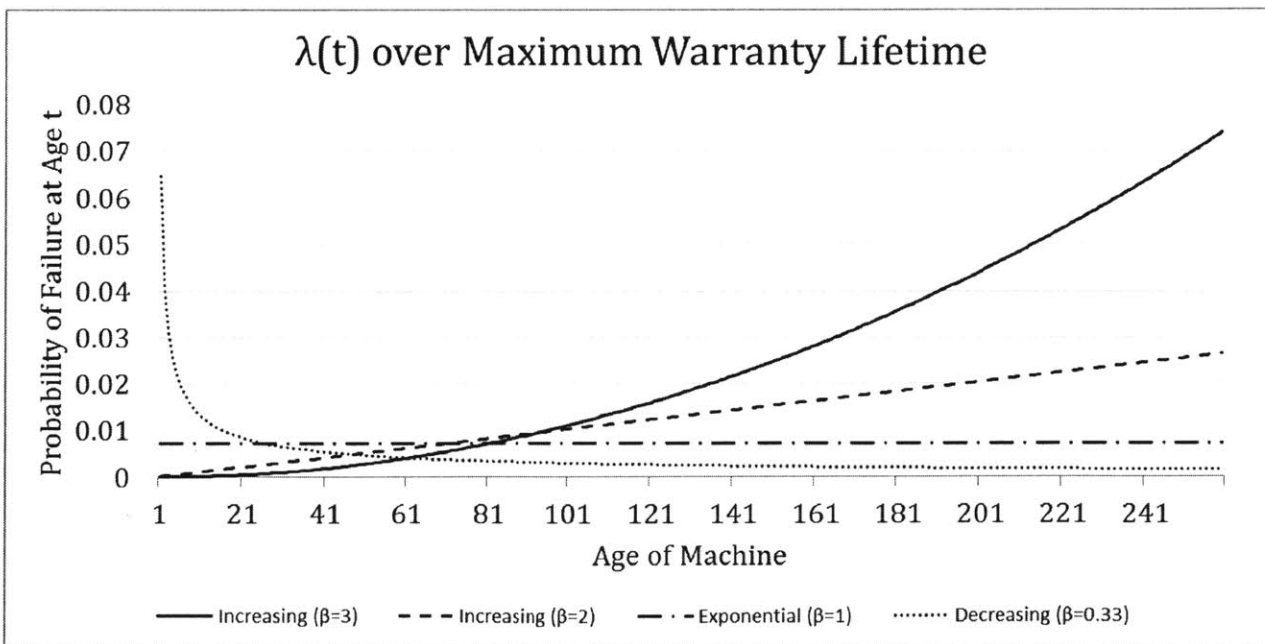


Figure 3.3.2. Weibull failure rate function at $\beta=0.33, \beta=1, \beta=2$ & $\beta=3$.

By generating random variates from the Weibull distribution with $\beta=1$, failures in each time period are simulated. We assume that a repaired machine is “like new”, so if a

machine had failed in period t and was repaired, in period $t+1$ we reset its age to 1. By comparing the probability of failure for each machine of age a_t in each time period to a random number generated at the time of machine instantiation or repair, we simulate failures in each period t . Because a machine failure creates demand for a spare part, the sum of all expected failures in each period becomes the demand curve for spare parts planning. An example from a single simulation run of the number of failures for each period is shown below in Figure 3.3.3. The time horizon is one year and each period is one day.




Figure 3.3.3. Simulated demand curve based off of expected failures.

3.4. Time-series Forecasting Methods

In our time-series forecasting model, the only data we take into account to forecast future spare parts demand is the historical demand data. Within this method, we test two

different exponential smoothing models; the simple exponential smoothing model and simple exponential smoothing model with trend, as proposed by Holt (1959), Winter (1960), and Brown (1959). Table 3.4.1 introduces the necessary notation to define these models.

Table 3.4.1. Time-series forecasting parameters

Notation	Description	Assigned Value
α	Alpha smoothing constant for level forecast	0.2
d_t	Demand in time t	given by failures/Bass model
$S_{t,t+1}$	Level forecast for time t+1 in time t	
$T_{t,t+1}$	Trend forecast for time t+1 in time t	
β	Beta smoothing constant for trend forecast	0.5

We choose an alpha (α) value of 0.2 and a beta (β) value of 0.5 based on values commonly used in industry. We initialize each model in time period 1 by setting both $F_{1,2}$ and $Ft_{1,2}$ to d_1 . The simple exponential smoothing model is given by:

$$F_{t,t+1} = \alpha * d_t + (1 - \alpha) * F_{t-1,t} \quad t = 1, \dots, n \quad (9)$$

The simple exponential smoothing model with trend is given by:

$$Ft_{t,t+1} = S_{t,t+1} + T_{t,t+1} \quad t = 1, \dots, n \quad (10)$$

where:

$$S_{t,t+1} = \alpha * d_t + (1 - \alpha) * (S_{t-1,t} + T_{t-1,t}) \quad t = 1, \dots, n \quad (11)$$

and where:

$$T_{t,t+1} = \beta * (F_{t,t+1} - F_{t-1,t}) + (1 - \beta) * (T_{t-1,t}) \quad t = 1, \dots, n \quad (12)$$

3.5. Machine Failure Prediction

Whereas in the time-series forecasting method we only use historical demand information to create our forecast, in the machine failure prediction method we assume that the machines in the installed base are generating some type of predictive signals, and that these signals can be received remotely through an Internet of Things infrastructure. These signals would be generated through the analysis of various machine diagnostic readings by sensors on each individual machine, such as temperature, vibration and/or the analysis of electronic machine logs, which contain output readings regarding the operating condition of the machine. In order to create a forecast using these signals, we utilize the binary classification method. A binary classification matrix, also known as a confusion matrix, is shown in Figure 3.5.1.

Table 3.5.1. Binary classification matrix.

	<i>Failures</i>	<i>No Failure</i>
<i>Signals</i>	<i>True Positive</i>	<i>False Positive</i>
<i>No Signals</i>	<i>False Negative</i>	<i>True Negative</i>

As can be seen in Table 3.5.1, the binary classification matrix is a method to measure the correlation of signals with failures, and vice-versa. For example, a true positive indicates the failure of a machine in the installed base that is able to be successfully detected ahead of time with a signal. While arbitrarily assigned at consistent intervals in our simulation model, the values in this matrix would be based off of the results of an analysis of using a set of machines' operating data to effectively predict failure in advance.

We use this matrix to define two key measurements that our predictive forecast is built on. Table 3.5.2 provides the notation required to define these two equations.

Table 3.5.2. Binary classification parameters.

Notation	Description	Alternative name
F_s	Count of machine periods in predictive window with a failure and a signal	True positive (TP)
F_{ns}	Count of machine periods in predictive window with a failure and no signal	False negative (FN)
NF_s	Count of machine periods in predictive window with no failure and a signal	False positive (FP)
NF_{ns}	Count of machine periods in predictive window with no failure and no signal	True negative (TN)

Note here that a machine-period is defined as the existence of a unique machine during some time period. For example, based on the Bass Model results from Section 3.3, three machines were under warranty in time period one, while six machines were under warranty in time period two, creating nine machine periods in total through period two. In total, throughout the entire demand horizon of periods 1 through 378, we find there are 269,543 machine periods in the installed base we generate.

The first equation, True Positive Rate (TPR), is given by:

$$TPR = \frac{\sum_{t=1}^n F_{S_t}}{(\sum_{t=1}^n F_{S_t} + \sum_{t=1}^n F_{ns_t})} \quad (13)$$

The second equation, the True Negative Rate (TNR), is given by:

$$TNR = \frac{\sum_{t=1}^n NF_{ns_t}}{(\sum_{t=1}^n NF_{ns_t} + \sum_{t=1}^n NF_{s_t})} \quad (14)$$

The TPR is a measure of how many machine periods with a failure we are able to accurately predict with a signal ahead of time, while the TNR measures how many machine periods with a non-failure we are able to not predict with a signal ahead of time. Note here that in the case of assigning signals, we assume in our model that a signal seen in time period t predicts a failure in time period $t+L+R$. This assumption allows our model to gain the maximum possible benefit in a predictive analytics scenario, as we assume a signal alerts us to a pending failure with enough time to dispatch a spare part and have it arrive in the same time period the failure occurs (the time to failure is greater than the total lead time to receive the spare part to fix the failure). If the period we are able to predict failures through is less than $t+L+R$, we would only be able to reduce the period of time the machine is failed, and not fix it proactively. This scenario still has benefits, as it allows us to reduce the potential stock out penalty paid by a company (by reducing the total downtime of a machine), but does not allow them to capture the full benefit of predicting failure $t+L+R$ periods ahead of time. We also note here that our model runs on a lead time of 1 and a review period of 1, creating a prediction window of 2. We did not test the efficacy of our model at longer combined lead times and review periods, and discuss that limitation in Section 5.

In addition to the TPR and the TNR, we also define the Positive Predictive Power (PPV) and Negative Predictive Power (NPV) of our model. The PPV is given by:

$$PPV = \frac{\sum_{t=1}^n F_{s_t}}{(\sum_{t=1}^n F_{s_t} + \sum_{t=1}^n NF_{s_t})} \quad (15)$$

Additionally, the NPV is given by:

$$NPV = \frac{\sum_{t=1}^n NF_{ns_t}}{(\sum_{t=1}^n NF_{ns_t} + \sum_{t=1}^n F_{ns_t})} \quad (16)$$

Our predictive forecast model works by allowing the user of the model to assign a TPR and TNR value. We then use these values to assign the appropriate amount of signals to create a forecast with. In order to do this, our simulation model first generates a number of failures using the logic displayed in Section 3.3. This step is taken for each individual simulation run. We then sum up the total amount of machine-periods with and without a failure throughout the entire demand horizon. While the total amount of machine-periods in our forecast is fixed for all simulation runs at 269,543, the number of failures that occurs during any distinct simulation run varies. We then use the TPR and TNR values to assign signals to machine periods. For example, by specifying TPR and TNR values of 0.9 and 0.9, respectively, we randomly assign signals to 90% of the machine periods in our forecast with a failure, and to 10% (1-TNR, or false positive rate (FPR)) of the machine periods in our forecast without a failure. The necessary notation for the equations that define the signals, non-signals, and forecast for the failure prediction model are provided in Table 3.5.3.

Table 3.5.3. Predictive forecast parameters.

Notation	Description	Assigned Value
P	prediction window	$L+R = 2$
S_j	1 if failure signal from machine j in time t , 0 otherwise	based on user input
NS_j	0 if non-failure signal from machine j in time t , 1 otherwise	based on user input
IB_t	size of the installed base at time t	based on Bass model results

The number of signals in each period is given by:

$$X_{t,t+P} = \sum_{j=1}^{IB_t} S_{j_t} \quad t = 1, \dots, n \quad (17)$$

Alternatively, the number of non-signals in each period is given by:

$$Y_{t,t+P} = \sum_{j=1}^{IB_t} NS_{j_t} \quad t = 1, \dots, n \quad (18)$$

We then use these signals to create a forecast for each time period in the demand horizon. This step considers the signals and non-signals assigned to each machine period using equations (17) and (18), and determines their accuracy by using the derived PPV and NPV from the confusion matrix as defined in (15) and (16). Our forecast for each period becomes:

$$F_{t,t+P} = X_{t,t+P} * PPV + Y_{t,t+P} * (1 - NPV) \quad t = 1, \dots, n \quad (19)$$

Then, we evaluate the variance of the signals, assuming a binomial distribution.

Accordingly, the variance for the machine signals is given by:

$$V_{t,t+P} = X_{t,t+P} * PPV * (1 - PPV) + Y_{t,t+P} * NPV * (1 - NPV) \quad t = 1, \dots, n \quad (20)$$

However, because the installed base is constantly changing in size between new machines being sold and old machines falling out of their warranty period, our forecast for failures is not comprised solely of signals generated from machines. This is because the failures in time $t+P$, or $t+L+R$, could also be comprised of failures in machines that were introduced into the installed based between time t and $t+L+R$, and thus were not able to have any signal generated in time t . Note here that equations (21) – (24) are found to be valid based on a lead time and review period of one period each, and were not tested in a more generalizable format. Additionally, note that for equations (21) and (22), we use an estimate of the size of the installed base in time $t+1$, given that we don't know that information in time t . Accordingly, the forecast for this subset of the machines in period $t+P$ is given by:

$$C_{t,t+2} = (IB_{t+1} - IB_t) * \lambda \quad t = 1, \dots, n \quad (21)$$

Similar to the predictive forecast, the variance for this set of machines is given by:

$$D_{t,t+2} = (IB_{t+1} - IB_t) * \lambda * (1 - \lambda) \quad t = 1, \dots, n \quad (22)$$

Finally, we derive the total forecast and variance by adding the two forecast and variance statements together, respectively. Accordingly, the total forecast and variance are given by:

$$U_{t,t+2} = F_{t,t+2} + C_{t,t+2} \quad t = 1, \dots, n \quad (23)$$

$$W_{t,t+2} = V_{t,t+2} + D_{t,t+2} \quad t = 1, \dots, n \quad (24)$$

3.6. General Model Design

As observed in both Schradly (1967) and Allen and D'Esopo (1968), we assume our spare parts inventory network consists of a single location, which we refer to as the central distribution center, or CDC. In our simulation model, we assume this location performs all inventory procurement and dispatch operations. However, in order to keep our model as simple as possible and focus on the core question of the value of predictive analytics on forecasting and inventory planning, we choose to leave out inventory repair operations from our model. We discuss this decision in Section 6 as an area for future extensions of our work.

In each forecasting model, we assume all demand is fulfilled from the CDC. If the CDC has inventory in stock, it is dispatched to fulfill the demand. If the total amount of demand in a given period exceeds the amount of inventory on hand entering the period, any inventory on hand is dispatched, the inventory level exiting the period is recorded as zero and a stock out is recorded for that period, with no demand being backlogged. In addition to a stock out being recorded, we also count how many pieces of inventory short our model is to fulfilling all demand in periods where there was a stock out. All orders from the CDC to the vendor are assumed to have a deterministic lead time.

3.7. Time-series Model Design

The CDC inventory ordering process for the two time-series based models we developed is based on an R, S periodic review policy. Such a policy assumes that the facility reviews its inventory after each time period (periodically) and if that inventory level is under some level S, places an order of size S less the current inventory position. The equations outlining the concepts our model runs on are outlined in (25) - (35).

We define the reorder point (S_t) as the mean demand over lead time plus the review period (μ_{L+R}), plus some z-score multiplied by the root mean square error over lead time plus the review period ($RMSE_{L+R}$). The $RMSE_{L+R}$ replaces the more commonly used standard deviation outlined in most literature. However, because at some time t we do not know the true standard deviation of the demand, we resort to using the RMSE. Thus, the reorder point is given by:

$$S_t = \mu_{L+R_t} + z * RMSE_{L+R_t} \quad (25)$$

We also note that we round up S_t to the nearest integer value. We define μ as the sum of all demand over some number of the last n periods divided by n. In our model, a user can define and control n, in that the mean demand can be either the mean over the entire planning horizon, or the mean over some rolling subset of the horizon (e.g., the last 10 periods of demand only). μ_{L+R_t} is given by:

$$\mu_{L+R_t} = \frac{\sum_{i=0}^{n-1} d_{t-i}}{n} * (L + R) \quad (26)$$

The $RMSE_{L+R_t}$ is defined in a series of steps. We start with the squared error (SE_t), defining it as the demand observed in time t (d_t), less the forecast made for time t in time $t-1$ ($F_{t,t-1}$), squared. SE_t is given by:

$$SE_t = (d_t - F_{t-1,t})^2 \quad (27)$$

The next step is to convert the SE_t into the mean squared error (MSE_t). As shown with μ , this is equal to the sum of all SE_t over some number of n periods divided by n . Here also, we allow the user to control n . Thus, MSE_t is given by:

$$MSE_t = \frac{\sum_{i=0}^{n-1} SE_{t-i}}{n} \quad (28)$$

To convert the MSE_t into the mean square error over lead time (MSE_{L+R_t}), we simply multiply MSE_t by the lead time plus the review period. Thus, MSE_{L+R_t} is given by:

$$MSE_{L+R_t} = MSE_t * (L + R) \quad (29)$$

Finally, we define the $RMSE_{L+R_t}$ as the square root of MSE_{L+R_t} :

$$RMSE_{L+R_t} = \sqrt{MSE_{L+R_t}} \quad (30)$$

The z-score is simply defined by the level of service we desire our model to support. By using a z-score, we are implicitly assuming that the errors are independently and identically (IID) distributed, and follow a normal distribution. A 50% CSL would be equivalent to a z-score of 0, a 95% CSL would be equivalent to a z-score of 1.645, a 99% CSL would be equivalent to a z-score of 2.33, and so on.

Next, we define our quantity purchased from the vendor (Qp_t). We define Qp_t as $S_t - Ip_t$. Thus, Qp_t is given by:

$$Qp_t = \max (S_t - IP_t, 0) \quad (31)$$

Accordingly, we only place an order for Qp_t if our inventory position (IP_t) is less than S_t . We define IP_t as the sum of inventory on hand coming out of a period ($IOHout_t$) plus the orders in transit ($Otrans_t$). Thus, IP_t is given by:

$$IP_t = IOHout_t + Otrans_t \quad (32)$$

We define $IOHout_t$ as the inventory on hand coming into a period ($IOHin_t$) less d_t . Thus, $IOHout_t$ is given by:

$$IOHout_t = \max (IOHin_t - d_t, 0) \quad (33)$$

Accordingly, $IOHin_t$ is given by the sum of $IOHout_{t-1}$ and the orders received from the vendor in the previous period (Or_{t-1}). Thus, $IOHin_t$ is given by:

$$IOHin_t = IOHout_{t-1} + Or_{t-1} \quad (34)$$

Or_t is simply Qp_t placed L periods ago. Thus, Or_t is given by:

$$Or_t = Qp_{t-L} \quad (35)$$

Finally, we define the $Otrans_t$ as the sum of all Qp placed within the last L periods.

Thus, $Otrans_t$ is defined as:

$$Otrans_t = \sum_{i=1}^{i=L} Qp_{t-i} \quad (36)$$

3.8. Predictive Forecast Model Design

The CDC ordering process for the predictive forecast model works under different assumptions. Equations (31)-(36) defined in Section 3.7 remain valid for the predictive forecast model. However, the reorder point, S_t , for the predictive forecast model is based on equations (23) and (24) from Section 3.5. Similar to those equations, (37) is defined based on a lead time and review period of one period each. Accordingly, the reorder point for the predictive forecast model is defined as:

$$S_t = U_{t-1,t+1} + U_{t,t+2} + z * (\sqrt{W_{t-1,t+1} + W_{t,t+2}}) \quad (37)$$

The first term in (37) is the predicted demand in each period between time t and $t+L+R$, while the second term is the standard deviation of the predictive forecast multiplied by some level of service, z . Figure 3.8.1 shows the period of time that these terms are designed to cover. Again, it is important to emphasize here that this approach was only tested at a lead time and review period of 1 period each, respectively.




Figure 3.8.1. Periods of demand covered by R, S periodic review policy

We note here that the reorder point is also rounded up to the next closest integer as was done in the time-series forecasting model. Additionally, the ordering logic for the predictive forecast model remains the same as in the time-series forecasting model, with an order placed anytime the inventory position falls at or below the reorder point.

3.9. Simulation Approach

In the first set of simulations, we generate a baseline set of summary statistics using the two time-series forecasting models, running each of the two models fifteen times each to generate a sample of performance characteristics. We then analyze which time-series model(s) provides the best performance, and select that model to compare the performance of the predictive analytics model against. Next, we run our predictive failure

forecasting model fifteen times each at a variety of different TPR and FPR levels, to measure the robustness of the model under different prediction levels. We summarize the results of these simulations in Section 4 of this thesis, Results and Data Analysis. For each distinct simulation, new failures (demand) are generated, while for each parameter setting within a simulation, new signals (predictions) are generated. Each predictive forecast simulation is run fifteen times at each unique combination of TPR and FPR in increments of 10% each from 0 to 1, keeping the same lead time and review period (2) and level of service (95%) as were used in the time-series forecasting model. Additionally, we also note here that each simulation for all three models tested was initially stocked with 10 units of inventory in time period 1 in order to cover the initial periods of demand. Finally, we also note that the number of periods, n , that was used for equations (26) and (28) was 10.

We also test our model at a much tighter range of specificity settings that stay within a closer range of 1.0. In this set of simulations, the TPR is still run in 10% increments, while the FPR is run at 0.1% increments. Again, each simulation is run fifteen times, with the TPR varying from 1 to 0, while the FPR varies from 1 to 0.99. We explain the rationale of this secondary set of testing in Section 4.

4. Results & Data Analysis

This section provides the quantitative results of the analyses conducted to determine the relationship between predictor accuracy and average inventory level, CSL, and IFR. First, we present a variety of summary statistics from each of the simulation runs described in Section 3.9 in table form. Next, we also present those same statistics in visual form via the use of bar charts and scatterplots.

4.1 Summary Statistics

First, we present the findings from the two time-series forecasting models. Table 4.1.2 shows the output metrics for each model in terms of both the average and standard deviation for each metric. Each metric is based off of the aggregate data from each of the fifteen simulation runs. Additionally, we also list the standard error of the standard deviations of the sample mean for the average inventory, CSL and IFR metrics. The notation required to define the standard error is provided in Table 4.1.1.

Table 4.1.1. Standard error parameters.

Notation	Description
s	standard deviation of sample mean
n	number of samples (simulations)

The standard error is defined as:

$$SE = \frac{s}{\sqrt{n}} \quad (38)$$

The standard error provides us with a confidence level for the standard deviation of the sample mean.

Table 4.1.2. Time-series performance metrics.

Metric	Measurement	Simple Exponential Smoothing	Simple Exponential Smoothing with Trend
Avg. Demand	Average	4.907	4.907
	Standard Deviation	0.065	0.065
Avg. Inventory	Average	8.884	8.592
	Standard Deviation	0.181	0.184
	Standard Error	0.047	0.047
CSL	Average	96.50%	95.98%
	Standard Deviation	0.08%	0.09%
	Standard Error	0.0002	0.0002
IFR	Average	98.49%	98.26%
	Standard Deviation	0.04%	0.04%
	Standard Error	0.0001	0.0001
Avg. Error	Average	0.000	0.000
	Standard Deviation	0.003	0.002
Avg. Bias	Average	0.008	0.027
	Standard Deviation	1.068	0.800

Based on the results highlighted in Table 4.1.2, we can see the maximum standard error value is equal to 0.047. We choose to use the simple exponential smoothing with trend time-series forecasting model as a baseline for comparison to the predictive failure forecasting model. We choose this model due to its satisfactory service performance and lower average inventory when compared to the simple exponential smoothing model.

4.2. Summary Statistics for Predictive Failure Model

In the Appendix (Section 7), we highlight the same performance metrics for the predictive forecast model as were shown in Table 4.1.2 for the time-series forecasting models. We list out these metrics separately for each TPR and TPR combination that were

tested. In general, we observe two major trends. The first trend is that as either the TPR moves closer to 1 or the FPR moves closer to 0, the predictive forecasting model is able to hold less average inventory while providing a higher level of service in comparison to the simple exponential smoothing with trend forecasting model. The second trend we observe is that no benefit is provided by the predictive forecasting model when the TPR is equal to the FPR.

4.3. Comparison of Performance Between Models

Figure 4.3.1 presents the observed percent inventory reduction between the average inventory held during the time-series model simulations and the average inventory held during the predictive forecast model simulations. It also shows the CSL and IFR performance at TPR values from 0.0 to 1.0. The FPR in the figure remains at a constant value of 0.0. We present the necessary notation in Table 4.3.1 for the percent reduction in inventory (PRI) calculation that is shown in both Figures 4.3.1 and 4.3.2.

Table 4.3.1. Percent reduction in inventory parameters.

Notation	Description
$Avg. Inv_{ts}$	average inventory from time-series forecasting model
$Avg. Inv_{pf}$	average inventory from predictive forecast model

We define the percent reduction in inventory as:

$$PRI = \frac{Avg. Inv_{ts} - Avg. Inv_{pf}}{Avg. Inv_{ts}} * 100\% \quad (39)$$




Figure 4.3.1. Percent reduction in inventory at varying TPR.

Similarly, we also test this metric at varying FPR and holding TPR constant at 1.0.

See Figure 4.3.2.




Figure 4.3.2. Percent reduction in inventory at varying FPR.

Due to the large jump in PRI, CSL and IFR between FPR values of 0.1 and 0.0, we also decided to simulate the FPR in smaller increments of 0.01% between FPR values of 0.1 and 0.0. The results of that set of simulations are shown in Figure 4.3.3.

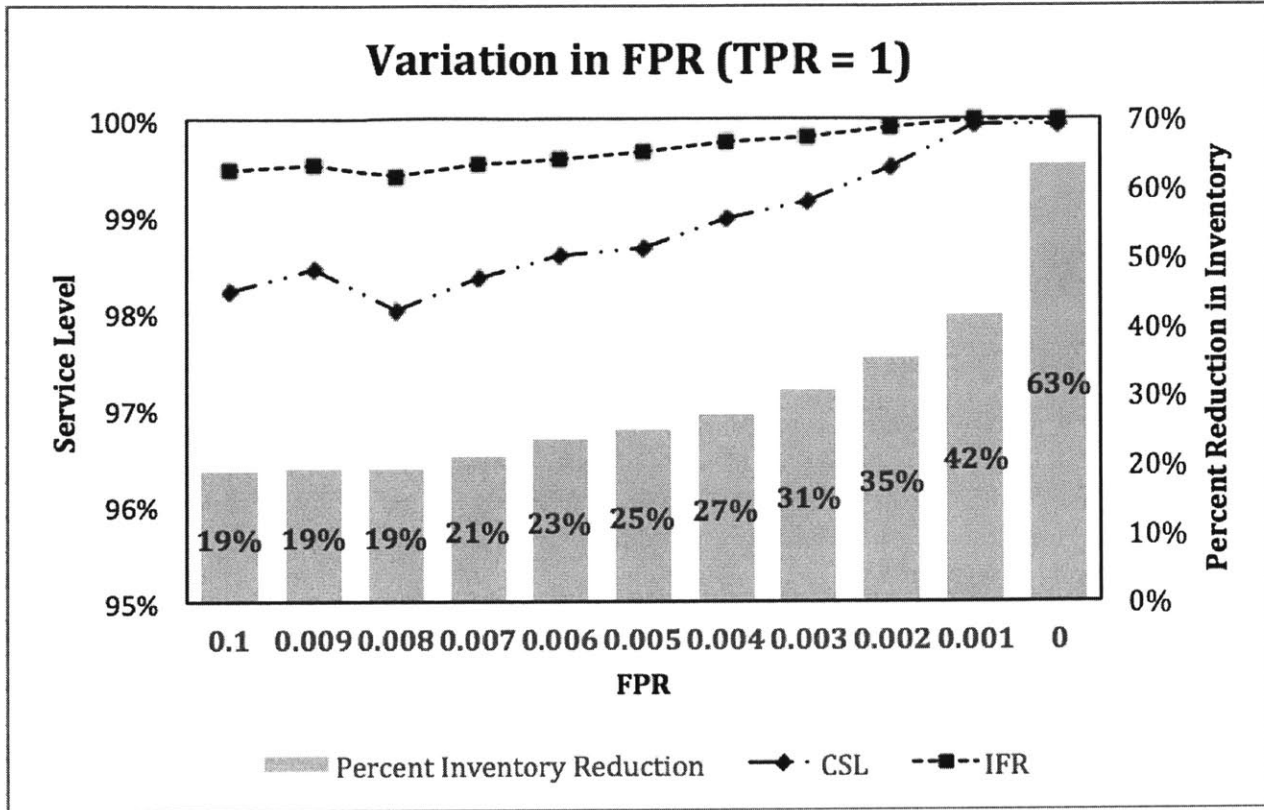


Figure 4.3.3. Percent reduction in inventory at smaller increments of varying FPR.

In Figure 4.3.3, we can see a much more gradual change in the PRI. In addition to Figures 4.3.1 – 4.3.3, Figure 4.3.4 gives all model results for combinations of TPR and FPR values between 0.0 and 1.0. Additionally, Figure 4.3.5 gives all model results for all combinations of TPR values between 0.0 and 1.0 and FPR values between 0.0 and 0.1.

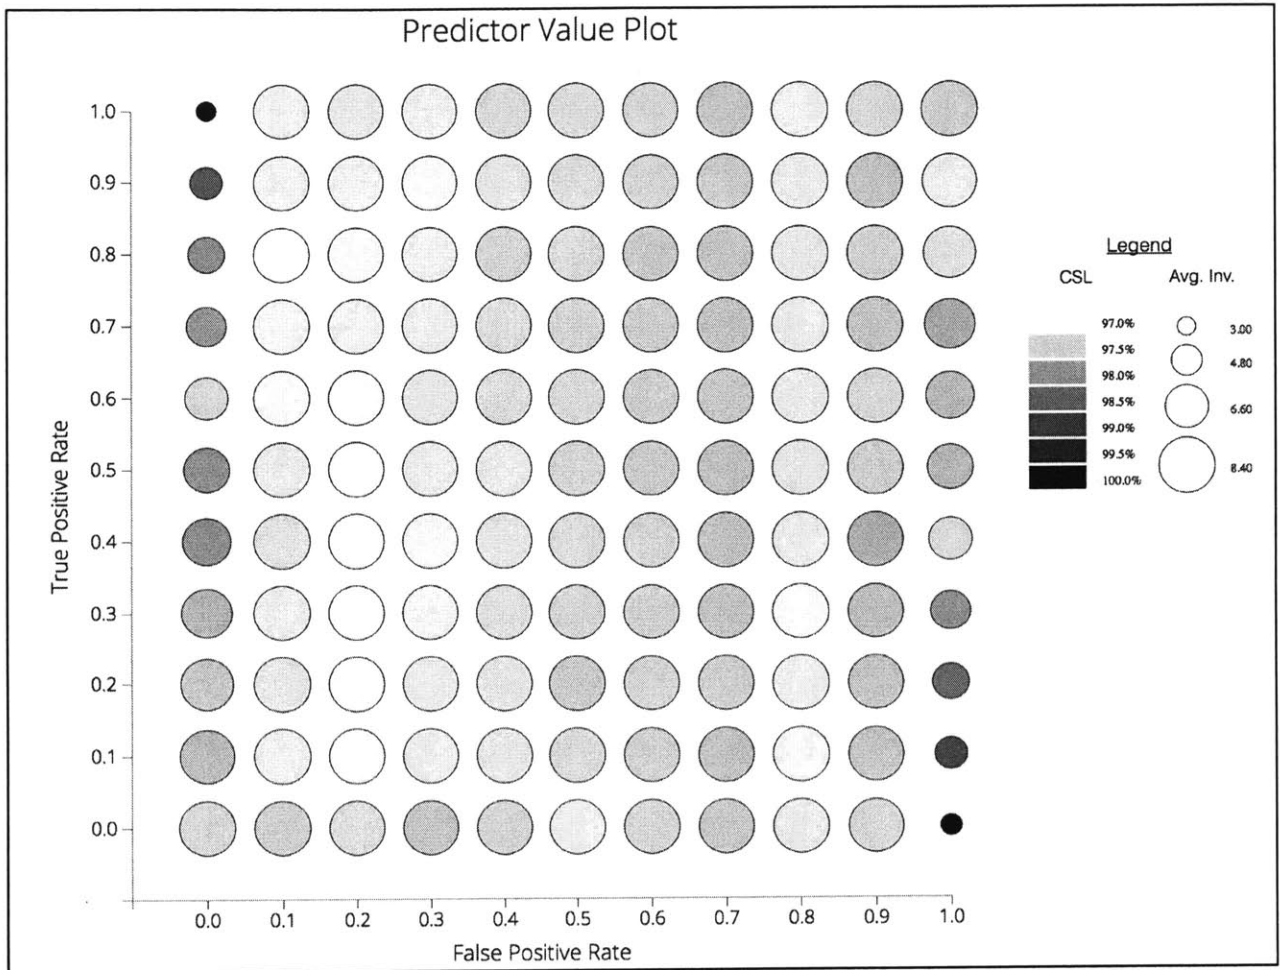


Figure 4.3.4. Plot of all combinations of TPR and FPR at 10% increments.




Figure 4.3.5. Plot of all combinations of TPR and FPR at 0.1% FPR intervals and 10% TPR intervals.

As we can see in Figures 4.3.4 and 4.3.5, the closer TPR is to 1 and/or the closer FPR is to 0, the lower the average inventory that can be achieved while maintaining a similar level of service. However, as discussed in Section 4.2, this is dependent on the TPR and FPR not being equal to one another. For example, at FPR = 1.0 and TPR = 1.0, as shown in the top right circle in Figure 4.3.4, the model has negligible benefits over the time-series forecasting model it was compared to. The farther away either the TPR or FPR value gets from its equivalent FPR or TPR value, respectively, the greater the benefit the model has. This can be observed in Figure 4.3.4 in the left and right most columns of circles.

5. Discussion

This section interprets the results discussed through Section 4. In addition to summarizing our findings and the general relationships we find in the data, we also explore the broader implications of our research and how they might be applied in a variety of service supply chain settings. We also discuss here the limitations and possible concerns we have based on the results we found.

5.1 Summary of Outcomes

After running the simulation model and analyzing the results, we find two general relationships among the inputs to the predictive forecasting model and its performance. First, we observe that the higher the rate of TPR or FPR (for a given value of the other), the greater the percent reduction in average inventory level, and the greater the percent improvement in CSL and IFR. This is able to be observed in Section 4 in Figures 4.3.2 – 4.3.5. Even at relatively low levels of TPR and FPR, a reduction in inventory is still observed, while holding a level of service at or slightly above the level of service obtained with the time-series forecasting model. This suggests a company does not need to have a highly predictive set of machine signals to see benefits from such a model.

The second general relationship we observe is that the model has value at all levels of TPR and FPR as long as the two values are not equal to one another. Again, this is able to be observed in Section 4 in Figure 4.3.4. These values being equal means a signal is just as likely to be a true positive as a false positive, and thus provides our model with no additional information with which to predict failure ahead of time.

One other relationship we observe, although are not able to prove based on the set of data analyzed, is that reductions in the FPR have a much greater effect on overall prediction performance than reductions in TPR. We believe this is due to the prevalence of failures in the data. We find that on average, for each simulation run our model generated 1,854.8 pieces of demand over the entire horizon. When considering that this demand is taking place over 269,543 machine periods, we find that the prevalence of failure is equal to 0.0069, or 0.69%. We believe that this explains the difference in reaction between changing TPR versus FPR in the model. For every 10% increase in the FPR, $(269,543 - 1,855) * 0.1 = 26,769$ false positives are added to the model, on average. However, for every 10% increase in the TPR, only $1,855 * 0.1 = 185$ true positives are added to the model, on average. This suggests that, at the failure rate our model was analyzed with, a reduction in the FPR is worth $26,769 / 185 = 144.69x$ the value that an equivalent percentage increase in the TPR is worth. This suggests it is much more important to maintain a high true negative rate than a high true positive rate. However, because we do not analyze alternative rates of failure, we are not able to prove this last observation here.

The results show that any improvement in predictive power, driven by changes in the TPR and FPR, contributes to a decrease in the error between the demand the forecast, provided the TPR and FPR values are not equal to one another. This also drives a natural decrease in the standard deviation, which explains the reduction in the average inventory levels. The greater the PPV of the model, the greater the reduction in inventory.

6. Conclusion

This section summarizes the major findings of our analysis, implications of the research on service supply chain design, and also discusses future related area of research and extension to our work.

6.1 Major findings

We summarize the major results of our analysis in the following section. First, in an inventory planning system that utilizes predictive analytics, any increase in the TPR or decrease in the FPR presents a benefit to spare parts forecasting and control, provided the $TPR \neq FPR$. This is due to the decrease in the variance as measured between the forecast and the actual demand when compared to a traditional time-series forecasting model.

In general, the higher the TPR and/or the higher the FPR, the greater the benefit received in terms of lower average inventories and higher average service levels, when compared to a time-series forecasting model. At a perfect predictive power, or when the TPR is equal to one and the FPR is equal to zero, a company should be able to hold the theoretical minimum inventory, which is simply the average demand per period, also known as the cycle stock.

Finally, assuming in general a company maintains a low failure rate in its products' components, we believe, although are not able to definitively prove, a predictive forecast inventory planning model is much more sensitive to decreases in the FPR versus equivalent decreases in the TPR. This is due to the number of false positives that are introduced to the model for every percent interval of change in the FPR versus the number of true positives that are introduced for every equivalent interval of change in the TPR.

6.2 Implications on Service Supply Chain Operations

As shown in Figures 4.3.1 and 4.3.2, predictive failure capability at any level of TPR and FPR (provided they do not equal one another) allows companies to reduce their average inventory, while also maintaining a higher or equivalent overall service level, when compared to traditional R, S inventory system combined with time-series forecasting models. One potential implication of this capability would be the redesign of a service supply chain network. As discussed in Section 2, service supply chain networks are often designed with multiple echelons, with the lowest level of echelons (on the local or regional level) containing many individual fixed sites or field service trunk stocks. Provided a company is able to:

- predict failure ahead of or at the lead time it takes to ship a unit of stock from the next highest echelon level to the local level, and
- execute this predictive approach across a number of its spare parts,

a company theoretically should be able to aggregate inventory at a more centralized location and reduce its overall service supply chain inventory at any given time. As discussed in Section 2.6.3, ideally a company would be able to use this approach on its category 'A' items in order to have the greatest effect on its total inventory investment. Due to the economics of the infrastructure required to take a predictive analytics approach, we anticipate that category 'A' items are where a predictive analytics naturally makes the most sense and where a company would want to install such a capability first. In the extreme best case scenario, not only would a company be able to reduce its total inventory investment and thus holding cost and obsolescence risk at any given time, it might also be able to remove an entire echelon of its service supply chain network. The predictive analytics

approach decouples the service supply chain from the “hierarchy of geography” as highlighted by Cohen et al. in Section 2.3.

Assuming a company is able to do this effectively on some certain set of its spare parts SKUs, the total inventory level for that set of spare parts should be able to be reduced by a factor of the \sqrt{n} , where n is the local number of facilities that a company is able to aggregate inventory across. This is due to the law of the sums of variances in the demand variability among all the different lowest-echelon sites. A more centralized inventory holding facility, beyond reducing total inventory levels, should also provide better control, cost and customer service by reducing phantom inventory levels, inventory shrinkage and obsolescence, and providing an overall lower stock out rate. In summary, predictive failure capability could drastically reduce the aggregate amount of inventory across a service supply chain network, as well as manage a greater amount of that inventory from a centralized location.

Predictive capability could also provide advantages to a company even when that capability is not able to extend beyond the lead time of shipment or repair. For example, as highlighted in Section 2, it is common in enterprise service supply chain operations to charge a stock out penalty or penalty cost per time unit of downtime for every unit of time a customer’s system or machine is down. While being able to predict failure outside of lead and repair time would be the ideal case, it would still be advantageous to service supply chain providers to predict failure some period in advance to reduce the total amount of downtime, and thus the total penalty cost paid. In essence, advance notice of failure would let service supply chain companies get a ‘jump-start’ on repair response, even if they weren’t able to completely prevent system failure of a machine.

Finally, we also note that perfect predictive power is not required to significantly increase the performance, both from a cost and service level, of a service supply chain's operations. One potential approach for a company to take with a PPV < 1 is to aggregate failure probabilities across a sub-section of the installed base assigned to a local depot. For example, if a company has five machines in its installed base serviced by a specific local depot, and each of those machines had generated a failure signal with a PPV of .5, a company should be able to use this information and proactively send three ($5 * .5 = 2.5 \rightarrow 3$) spare parts to the depot in order to service those repairs ahead of time. This is similar to the approach that Louit, Pascual, Banjevic and Jardine (2011) take in their work as highlighted in Section 2.

6.3 Extensions to our work and avenues for further exploration

There are still many other further areas for exploration in this field. One such area is an extension of the lead and review periods used in this model. Periods of one and one were chosen respectively in order to keep the analysis as simple as possible. However, in practice, due to financial and general operational overhead constraints, it is much more common to find lead times longer than one period, as well as to find companies who review their inventory position on a longer interval. Accordingly, an interesting extension of the work we propose in this thesis is to modify our model to work in a lead time plus review period longer than two. We believe a longer lead time plus review period would provide even more significant benefits than the benefits found in this paper, due to larger amounts of safety stock needing to be held to cover the longer periods between order placement and order arrival.

The impact of repair operations on the inventory model developed could also be explored. The additional source of inventory that comes from repaired parts adds complexity to the model. A new inventory policy will need to be developed that should take into account the timing of repair arrival, the repair rate and the repair lead time. We believe analyzing the effects that varying repair rates has on the performance of the model would be an interesting area for further research.

It would also be interesting to determine a method for deciding which types of products can be economically monitored and predictive analytics applied. As discussed in Section 2, OEMs typically need to contend with warranty requests over a much longer span of time than a spare part needed to fill those requests might be available for (the “lifecycle mismatch” problem). In order to keep our analysis relatively simple, we chose to ignore that constraint; however, it is an interesting area for further research. Similarly, we also chose to ignore any lot sizing requirements or minimum ordering quantities (MOQs) with the vendor. In general, we expect that the higher the MOQ/batch size, the higher the average inventory; however, given that this inventory increase would take place both within the time-series and predictive analytics model, it is plausible that the net effects would be negligible.

One additional avenue for research would be an investigation of the economics of machine failure prediction and on what products and/or spare parts installing predictive analytics capabilities makes the most sense. As described above, our model is allowed to order components in any quantity in any time period. However, the expense of many single shipments, in addition to the overall cost of the infrastructure required to effectively predict failure, will more than likely outweigh the benefits for low cost, non-critical parts

for system performance. We anticipate where the Internet of Things has the most potential on inventory operations is on higher expense, system critical parts ('A' items) that represent a large investment of working capital on the part of the service supply chain.

Finally, one last modification of our work that we believe is a potential area for further research is modifying the Weibull parameters used to generate our spare part demand. Specifically, we would be interested in seeing research that tests the performance of our model with a β value greater than 1, which as discussed in section 3.3, represents an increasing rate of failure over time. We suspect there are many spare parts components that are increasingly likely to fail as they increase in age, and thus that our model needs to be modified in order to correctly predict failure on those parts.

In summary, there are many real world challenges that make implementing the approach taken in this research into a real world supply planning situation. However, we believe the approach outlined in this thesis begins to show the potential service supply chain cost savings that are possible with the implementation of the Internet of Things.

7. Appendix

7.1 Performance metrics for all combinations of TPR and FPR across fifteen simulations

Table 7.1.1 Average Inventory for all combinations of TPR and FPR

TPR	Values	FPR											Total
		0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1	
0	Average of Average Inventory	8.25	8.24	8.24	8.22	8.24	8.24	8.22	8.20	8.18	8.09	3.37	7.43
	StdDev of Average Inventory	0.11	0.11	0.11	0.10	0.11	0.11	0.12	0.11	0.12	0.11	0.07	1.79
	StdError of Average Inventory	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.02	0.46
0.1	Average of Average Inventory	8.03	8.29	8.30	8.24	8.29	8.28	8.17	8.23	8.25	8.14	4.95	7.94
	StdDev of Average Inventory	0.13	0.17	0.09	0.11	0.17	0.11	0.13	0.12	0.11	0.11	0.07	0.93
	StdError of Average Inventory	0.03	0.04	0.02	0.03	0.04	0.03	0.03	0.03	0.03	0.03	0.02	0.24
0.2	Average of Average Inventory	7.75	8.29	8.30	8.24	8.29	8.29	8.17	8.24	8.26	8.15	5.53	7.97
	StdDev of Average Inventory	0.16	0.16	0.09	0.10	0.17	0.11	0.13	0.12	0.10	0.11	0.07	0.77
	StdError of Average Inventory	0.04	0.04	0.02	0.03	0.04	0.03	0.03	0.03	0.03	0.03	0.02	0.20
0.3	Average of Average Inventory	7.47	8.28	8.30	8.25	8.30	8.27	8.17	8.24	8.27	8.17	5.99	7.99
	StdDev of Average Inventory	0.14	0.15	0.10	0.11	0.17	0.12	0.12	0.12	0.12	0.10	0.06	0.66
	StdError of Average Inventory	0.04	0.04	0.03	0.03	0.04	0.03	0.03	0.03	0.03	0.02	0.02	0.17
0.4	Average of Average Inventory	7.17	8.27	8.29	8.24	8.30	8.28	8.17	8.25	8.27	8.19	6.47	8.00
	StdDev of Average Inventory	0.06	0.15	0.11	0.10	0.17	0.11	0.12	0.11	0.10	0.10	0.06	0.57
	StdError of Average Inventory	0.02	0.04	0.03	0.03	0.04	0.03	0.03	0.03	0.03	0.03	0.02	0.15
0.5	Average of Average Inventory	6.82	8.26	8.27	8.23	8.29	8.28	8.17	8.25	8.27	8.21	6.83	8.00
	StdDev of Average Inventory	0.08	0.15	0.12	0.11	0.17	0.11	0.13	0.11	0.11	0.10	0.07	0.56
	StdError of Average Inventory	0.02	0.04	0.03	0.03	0.04	0.03	0.03	0.03	0.03	0.03	0.02	0.14
0.6	Average of Average Inventory	6.45	8.23	8.27	8.23	8.28	8.27	8.17	8.24	8.26	8.21	7.18	7.98
	StdDev of Average Inventory	0.08	0.16	0.13	0.12	0.17	0.11	0.13	0.11	0.10	0.11	0.07	0.58
	StdError of Average Inventory	0.02	0.04	0.03	0.03	0.04	0.03	0.03	0.03	0.03	0.03	0.02	0.15
0.7	Average of Average Inventory	5.98	8.21	8.25	8.21	8.27	8.26	8.16	8.24	8.26	8.22	7.50	7.96
	StdDev of Average Inventory	0.08	0.15	0.12	0.13	0.17	0.11	0.13	0.10	0.10	0.11	0.13	0.68
	StdError of Average Inventory	0.02	0.04	0.03	0.03	0.04	0.03	0.03	0.03	0.03	0.03	0.03	0.17
0.8	Average of Average Inventory	5.50	8.18	8.23	8.20	8.25	8.26	8.15	8.24	8.24	8.21	7.75	7.93
	StdDev of Average Inventory	0.08	0.16	0.12	0.12	0.16	0.10	0.12	0.10	0.11	0.10	0.13	0.79
	StdError of Average Inventory	0.02	0.04	0.03	0.03	0.04	0.03	0.03	0.03	0.03	0.03	0.03	0.20
0.9	Average of Average Inventory	4.92	8.15	8.23	8.19	8.24	8.24	8.15	8.23	8.26	8.23	8.01	7.89
	StdDev of Average Inventory	0.08	0.15	0.12	0.12	0.17	0.11	0.13	0.10	0.10	0.11	0.11	0.96
	StdError of Average Inventory	0.02	0.04	0.03	0.03	0.04	0.03	0.03	0.03	0.03	0.03	0.03	0.25
1	Average of Average Inventory	3.14	8.11	8.19	8.17	8.23	8.24	8.14	8.22	8.25	8.21	8.28	7.74
	StdDev of Average Inventory	0.08	0.16	0.12	0.13	0.16	0.12	0.12	0.11	0.11	0.11	0.11	1.47
	StdError of Average Inventory	0.02	0.04	0.03	0.03	0.04	0.03	0.03	0.03	0.03	0.03	0.03	0.38
Total Average of Average Inventory		6.50	8.23	8.26	8.22	8.27	8.27	8.17	8.23	8.25	8.19	6.25	7.89
Total StdDev of Average Inventory		1.47	0.16	0.12	0.11	0.16	0.11	0.12	0.11	0.11	0.11	1.63	0.99
Total StdError of Average Inventory		1.47	0.16	0.12	0.11	0.16	0.11	0.12	0.11	0.11	0.11	1.63	0.99

Table 7.1.2 Cycle Service Level for all combinations of TPR and FPR

TPR	Values	FPR										Total	
		0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9		1
0	Average of CSL	97.54%	97.60%	97.51%	97.65%	97.53%	97.32%	97.53%	97.61%	97.37%	97.53%	100.00%	97.92%
	StdDev of CSL	0.52%	0.47%	0.56%	0.58%	0.56%	0.57%	0.46%	0.44%	0.73%	0.62%	0.00%	1.05%
	StdError of CSL	0.13%	0.12%	0.14%	0.15%	0.14%	0.15%	0.12%	0.11%	0.19%	0.16%	0.00%	0.27%
0.1	Average of CSL	97.73%	97.37%	97.22%	97.36%	97.41%	97.53%	97.58%	97.68%	97.27%	97.66%	98.89%	97.60%
	StdDev of CSL	0.71%	0.71%	0.55%	0.79%	0.80%	0.42%	0.76%	0.68%	0.71%	0.65%	0.54%	0.78%
	StdError of CSL	0.18%	0.18%	0.14%	0.20%	0.21%	0.11%	0.20%	0.18%	0.18%	0.17%	0.14%	0.20%
0.2	Average of CSL	97.63%	97.41%	97.19%	97.36%	97.44%	97.61%	97.54%	97.56%	97.32%	97.63%	98.45%	97.55%
	StdDev of CSL	0.55%	0.67%	0.56%	0.82%	0.78%	0.39%	0.80%	0.71%	0.69%	0.77%	0.56%	0.72%
	StdError of CSL	0.14%	0.17%	0.15%	0.21%	0.20%	0.10%	0.21%	0.18%	0.18%	0.20%	0.15%	0.19%
0.3	Average of CSL	97.77%	97.37%	97.12%	97.31%	97.48%	97.54%	97.56%	97.66%	97.24%	97.68%	98.05%	97.52%
	StdDev of CSL	0.50%	0.78%	0.64%	0.83%	0.81%	0.36%	0.68%	0.67%	0.66%	0.90%	0.81%	0.73%
	StdError of CSL	0.13%	0.20%	0.17%	0.21%	0.21%	0.09%	0.17%	0.17%	0.17%	0.23%	0.21%	0.19%
0.4	Average of CSL	98.07%	97.48%	97.19%	97.29%	97.46%	97.53%	97.56%	97.70%	97.37%	97.85%	97.53%	97.55%
	StdDev of CSL	0.65%	0.81%	0.51%	0.80%	0.81%	0.39%	0.72%	0.63%	0.64%	0.70%	0.51%	0.69%
	StdError of CSL	0.17%	0.21%	0.13%	0.21%	0.21%	0.10%	0.18%	0.16%	0.16%	0.18%	0.13%	0.18%
0.5	Average of CSL	98.04%	97.37%	97.10%	97.36%	97.39%	97.54%	97.60%	97.68%	97.39%	97.56%	97.77%	97.53%
	StdDev of CSL	0.60%	0.73%	0.62%	0.84%	0.76%	0.41%	0.74%	0.65%	0.73%	0.63%	0.75%	0.71%
	StdError of CSL	0.16%	0.19%	0.16%	0.22%	0.20%	0.11%	0.19%	0.17%	0.19%	0.16%	0.19%	0.18%
0.6	Average of CSL	97.54%	97.25%	97.20%	97.39%	97.46%	97.49%	97.58%	97.61%	97.34%	97.49%	97.72%	97.46%
	StdDev of CSL	0.76%	0.71%	0.52%	0.84%	0.82%	0.43%	0.75%	0.71%	0.56%	0.76%	0.48%	0.68%
	StdError of CSL	0.20%	0.18%	0.13%	0.22%	0.21%	0.11%	0.19%	0.18%	0.15%	0.20%	0.12%	0.17%
0.7	Average of CSL	98.02%	97.27%	97.31%	97.32%	97.44%	97.53%	97.56%	97.66%	97.32%	97.61%	97.83%	97.53%
	StdDev of CSL	0.65%	0.71%	0.60%	0.83%	0.72%	0.44%	0.74%	0.67%	0.61%	0.83%	0.53%	0.69%
	StdError of CSL	0.17%	0.18%	0.16%	0.21%	0.19%	0.11%	0.19%	0.17%	0.16%	0.21%	0.14%	0.18%
0.8	Average of CSL	98.12%	97.00%	97.24%	97.32%	97.53%	97.48%	97.60%	97.63%	97.36%	97.53%	97.39%	97.47%
	StdDev of CSL	0.73%	0.86%	0.50%	0.82%	0.80%	0.44%	0.73%	0.70%	0.64%	0.88%	0.72%	0.75%
	StdError of CSL	0.19%	0.22%	0.13%	0.21%	0.21%	0.11%	0.19%	0.18%	0.16%	0.23%	0.19%	0.19%
0.9	Average of CSL	98.72%	97.32%	97.31%	97.24%	97.39%	97.43%	97.51%	97.58%	97.34%	97.63%	97.31%	97.53%
	StdDev of CSL	0.52%	0.58%	0.75%	0.93%	0.71%	0.43%	0.75%	0.70%	0.71%	0.83%	0.54%	0.78%
	StdError of CSL	0.13%	0.15%	0.19%	0.24%	0.18%	0.11%	0.19%	0.18%	0.18%	0.22%	0.14%	0.20%
1	Average of CSL	99.93%	97.32%	97.37%	97.32%	97.48%	97.46%	97.51%	97.63%	97.31%	97.51%	97.55%	97.67%
	StdDev of CSL	0.12%	0.78%	0.64%	0.79%	0.70%	0.48%	0.78%	0.72%	0.60%	0.94%	0.69%	0.99%
	StdError of CSL	0.03%	0.20%	0.16%	0.20%	0.18%	0.12%	0.20%	0.19%	0.16%	0.24%	0.18%	0.25%
Total Average of CSL		98.10%	97.34%	97.25%	97.36%	97.45%	97.50%	97.56%	97.64%	97.33%	97.61%	98.22%	97.58%
Total StdDev of CSL		0.88%	0.71%	0.59%	0.79%	0.73%	0.43%	0.70%	0.65%	0.65%	0.76%	1.07%	0.80%
Total StdError of CSL		0.23%	0.18%	0.15%	0.20%	0.19%	0.11%	0.18%	0.17%	0.17%	0.20%	0.28%	0.21%

Table 7.1.3 Item Fill Rate for all combinations of TPR and FPR

TPR	Values	FPR											Total
		0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1	
0	Average of IFR	99.00%	99.00%	98.97%	99.00%	98.96%	98.92%	98.99%	98.98%	98.97%	99.02%	100.00%	99.15%
	StdDev of IFR	0.29%	0.27%	0.26%	0.27%	0.27%	0.26%	0.18%	0.34%	0.28%	0.29%	0.00%	0.45%
	StdError of IFR	0.08%	0.07%	0.07%	0.07%	0.07%	0.07%	0.05%	0.09%	0.07%	0.07%	0.00%	0.12%
0.1	Average of IFR	99.07%	98.94%	98.81%	98.90%	99.04%	98.96%	99.00%	99.06%	98.99%	99.05%	99.73%	99.05%
	StdDev of IFR	0.33%	0.29%	0.37%	0.32%	0.26%	0.22%	0.45%	0.25%	0.34%	0.36%	0.14%	0.38%
	StdError of IFR	0.08%	0.08%	0.10%	0.08%	0.07%	0.06%	0.12%	0.06%	0.09%	0.09%	0.04%	0.10%
0.2	Average of IFR	99.04%	98.95%	98.81%	98.87%	99.04%	98.99%	98.97%	99.03%	98.98%	99.07%	99.55%	99.02%
	StdDev of IFR	0.20%	0.30%	0.37%	0.35%	0.28%	0.22%	0.48%	0.27%	0.36%	0.39%	0.20%	0.36%
	StdError of IFR	0.05%	0.08%	0.09%	0.09%	0.07%	0.06%	0.12%	0.07%	0.09%	0.10%	0.05%	0.09%
0.3	Average of IFR	99.16%	98.92%	98.81%	98.87%	99.05%	98.97%	98.98%	99.04%	98.93%	99.02%	99.45%	99.02%
	StdDev of IFR	0.22%	0.32%	0.37%	0.35%	0.26%	0.23%	0.44%	0.30%	0.34%	0.40%	0.20%	0.35%
	StdError of IFR	0.06%	0.08%	0.09%	0.09%	0.07%	0.06%	0.11%	0.08%	0.09%	0.10%	0.05%	0.09%
0.4	Average of IFR	99.31%	99.00%	98.80%	98.87%	99.04%	98.97%	98.96%	99.03%	98.96%	99.04%	99.25%	99.02%
	StdDev of IFR	0.24%	0.31%	0.36%	0.35%	0.27%	0.23%	0.45%	0.26%	0.34%	0.39%	0.18%	0.34%
	StdError of IFR	0.06%	0.08%	0.09%	0.09%	0.07%	0.06%	0.12%	0.07%	0.09%	0.10%	0.05%	0.09%
0.5	Average of IFR	99.34%	98.96%	98.79%	98.88%	99.03%	98.96%	98.98%	99.05%	98.96%	98.98%	99.29%	99.02%
	StdDev of IFR	0.21%	0.31%	0.37%	0.34%	0.26%	0.23%	0.44%	0.25%	0.40%	0.40%	0.28%	0.35%
	StdError of IFR	0.05%	0.08%	0.10%	0.09%	0.07%	0.06%	0.11%	0.06%	0.10%	0.10%	0.07%	0.09%
0.6	Average of IFR	99.28%	98.92%	98.84%	98.92%	99.04%	98.95%	98.98%	99.04%	98.91%	98.97%	99.24%	99.01%
	StdDev of IFR	0.20%	0.36%	0.33%	0.38%	0.27%	0.22%	0.45%	0.28%	0.34%	0.44%	0.25%	0.35%
	StdError of IFR	0.05%	0.09%	0.08%	0.10%	0.07%	0.06%	0.12%	0.07%	0.09%	0.11%	0.07%	0.09%
0.7	Average of IFR	99.39%	98.93%	98.85%	98.87%	99.04%	98.96%	98.96%	99.05%	98.94%	98.99%	99.16%	99.01%
	StdDev of IFR	0.27%	0.31%	0.40%	0.34%	0.26%	0.23%	0.44%	0.26%	0.36%	0.44%	0.20%	0.35%
	StdError of IFR	0.07%	0.08%	0.10%	0.09%	0.07%	0.06%	0.11%	0.07%	0.09%	0.11%	0.05%	0.09%
0.8	Average of IFR	99.49%	98.92%	98.84%	98.89%	99.04%	98.95%	98.96%	99.03%	98.94%	98.96%	99.13%	99.01%
	StdDev of IFR	0.21%	0.35%	0.33%	0.32%	0.27%	0.23%	0.45%	0.27%	0.36%	0.44%	0.31%	0.36%
	StdError of IFR	0.05%	0.09%	0.09%	0.08%	0.07%	0.06%	0.12%	0.07%	0.09%	0.11%	0.08%	0.09%
0.9	Average of IFR	99.70%	98.88%	98.85%	98.88%	99.02%	98.98%	98.96%	99.02%	98.94%	98.98%	99.02%	99.02%
	StdDev of IFR	0.12%	0.34%	0.36%	0.32%	0.29%	0.23%	0.45%	0.30%	0.36%	0.43%	0.25%	0.39%
	StdError of IFR	0.03%	0.09%	0.09%	0.08%	0.07%	0.06%	0.12%	0.08%	0.09%	0.11%	0.06%	0.10%
1	Average of IFR	99.99%	98.92%	98.85%	98.87%	99.06%	98.97%	98.95%	99.02%	98.91%	98.96%	99.06%	99.05%
	StdDev of IFR	0.02%	0.28%	0.36%	0.32%	0.25%	0.22%	0.46%	0.26%	0.36%	0.44%	0.23%	0.43%
	StdError of IFR	0.01%	0.07%	0.09%	0.08%	0.07%	0.06%	0.12%	0.07%	0.09%	0.11%	0.06%	0.11%
Total Average of IFR		99.34%	98.94%	98.84%	98.89%	99.03%	98.96%	98.97%	99.03%	98.95%	99.00%	99.41%	99.03%
Total StdDev of IFR		0.36%	0.31%	0.35%	0.33%	0.26%	0.22%	0.42%	0.27%	0.34%	0.39%	0.39%	0.38%
Total StdError of IFR		0.36%	0.31%	0.35%	0.33%	0.26%	0.22%	0.42%	0.27%	0.34%	0.39%	0.39%	0.38%

Table 7.1.4 Average Error for all combinations of TPR and FPR

TPR	Values	FPR											Total
		0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1	
0	Average of Avg. Error	-0.0002	-0.0002	-0.0002	-0.0002	-0.0002	-0.0002	-0.0002	-0.0001	-0.0001	-0.0002	0.0000	-0.0001
	StdDev of Avg. Error	0.0000	0.0000	0.0000	0.0000	0.0000	0.0001	0.0001	0.0001	0.0001	0.0002	0.0000	0.0001
0.1	Average of Avg. Error	-0.0018	-0.0020	-0.0018	-0.0018	-0.0017	-0.0018	-0.0019	-0.0018	-0.0018	-0.0017	-0.0018	-0.0018
	StdDev of Avg. Error	0.0020	0.0002	0.0007	0.0007	0.0009	0.0007	0.0001	0.0007	0.0007	0.0009	0.0008	0.0009
0.2	Average of Avg. Error	-0.0023	-0.0038	-0.0037	-0.0037	-0.0035	-0.0036	-0.0037	-0.0036	-0.0037	-0.0034	-0.0037	-0.0035
	StdDev of Avg. Error	0.0024	0.0003	0.0007	0.0008	0.0009	0.0007	0.0001	0.0007	0.0007	0.0008	0.0008	0.0010
0.3	Average of Avg. Error	-0.0067	-0.0055	-0.0054	-0.0054	-0.0052	-0.0055	-0.0055	-0.0054	-0.0055	-0.0053	-0.0057	-0.0055
	StdDev of Avg. Error	0.0051	0.0006	0.0007	0.0007	0.0009	0.0007	0.0002	0.0007	0.0006	0.0009	0.0008	0.0017
0.4	Average of Avg. Error	-0.0087	-0.0071	-0.0075	-0.0072	-0.0071	-0.0073	-0.0073	-0.0072	-0.0073	-0.0070	-0.0073	-0.0074
	StdDev of Avg. Error	0.0070	0.0008	0.0009	0.0008	0.0009	0.0007	0.0003	0.0007	0.0007	0.0009	0.0010	0.0022
0.5	Average of Avg. Error	-0.0097	-0.0092	-0.0092	-0.0090	-0.0090	-0.0092	-0.0090	-0.0091	-0.0091	-0.0089	-0.0094	-0.0092
	StdDev of Avg. Error	0.0059	0.0006	0.0011	0.0009	0.0011	0.0008	0.0004	0.0007	0.0006	0.0009	0.0007	0.0019
0.6	Average of Avg. Error	-0.0116	-0.0118	-0.0108	-0.0111	-0.0107	-0.0109	-0.0109	-0.0109	-0.0110	-0.0106	-0.0108	-0.0110
	StdDev of Avg. Error	0.0077	0.0010	0.0009	0.0011	0.0009	0.0007	0.0005	0.0008	0.0007	0.0009	0.0009	0.0024
0.7	Average of Avg. Error	-0.0120	-0.0131	-0.0126	-0.0128	-0.0127	-0.0126	-0.0127	-0.0126	-0.0128	-0.0125	-0.0128	-0.0127
	StdDev of Avg. Error	0.0077	0.0016	0.0013	0.0011	0.0011	0.0008	0.0004	0.0007	0.0007	0.0010	0.0008	0.0024
0.8	Average of Avg. Error	-0.0163	-0.0151	-0.0148	-0.0147	-0.0144	-0.0145	-0.0144	-0.0145	-0.0145	-0.0143	-0.0146	-0.0147
	StdDev of Avg. Error	0.0088	0.0013	0.0012	0.0012	0.0011	0.0009	0.0005	0.0010	0.0007	0.0010	0.0012	0.0028
0.9	Average of Avg. Error	-0.0167	-0.0166	-0.0167	-0.0163	-0.0164	-0.0163	-0.0161	-0.0164	-0.0165	-0.0162	-0.0165	-0.0164
	StdDev of Avg. Error	0.0088	0.0013	0.0011	0.0011	0.0015	0.0009	0.0007	0.0008	0.0007	0.0010	0.0013	0.0028
1	Average of Avg. Error	-0.0188	-0.0182	-0.0184	-0.0185	-0.0180	-0.0178	-0.0181	-0.0183	-0.0182	-0.0180	-0.0183	-0.0182
	StdDev of Avg. Error	0.0096	0.0014	0.0012	0.0013	0.0013	0.0010	0.0007	0.0009	0.0007	0.0011	0.0008	0.0030
Total Average of Avg. Error		-0.0095	-0.0093	-0.0092	-0.0092	-0.0090	-0.0091	-0.0091	-0.0091	-0.0091	-0.0089	-0.0084	-0.0091
Total StdDev of Avg. Error		0.0088	0.0059	0.0059	0.0059	0.0058	0.0057	0.0057	0.0058	0.0058	0.0058	0.0062	0.0062

Table 7.1.5 Bias Error for all combinations of TPR and FPR

TPR	Values	FPR											Total
		0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1	
0	Average of Bias Error	-0.0639	-0.0681	-0.0642	-0.0635	-0.0721	-0.0659	-0.0682	-0.0464	-0.0555	-0.0646	0.0000	-0.0530
	StdDev of Bias Error	0.0015	0.0097	0.0089	0.0120	0.0164	0.0234	0.0262	0.0248	0.0306	0.0599	0.0000	0.0337
0.1	Average of Bias Error	-0.7178	-0.7727	-0.7116	-0.7064	-0.6580	-0.6869	-0.7600	-0.7127	-0.7150	-0.6540	-0.6946	-0.7082
	StdDev of Bias Error	0.7857	0.0773	0.2592	0.2659	0.3478	0.2556	0.0448	0.2650	0.2625	0.3329	0.3045	0.3341
0.2	Average of Bias Error	-0.9055	-1.4853	-1.4601	-1.4378	-1.3601	-1.3909	-1.4469	-1.4057	-1.4310	-1.3467	-1.4608	-1.3750
	StdDev of Bias Error	0.9543	0.1114	0.2727	0.3172	0.3606	0.2575	0.0544	0.2711	0.2696	0.3142	0.3160	0.4055
0.3	Average of Bias Error	-2.6251	-2.1530	-2.1031	-2.1068	-2.0153	-2.1312	-2.1642	-2.1146	-2.1399	-2.0542	-2.2263	-2.1663
	StdDev of Bias Error	2.0012	0.2226	0.2789	0.2717	0.3638	0.2720	0.0844	0.2830	0.2496	0.3424	0.3075	0.6584
0.4	Average of Bias Error	-3.4117	-2.7655	-2.9308	-2.8084	-2.7766	-2.8518	-2.8668	-2.8250	-2.8724	-2.7458	-2.8487	-2.8823
	StdDev of Bias Error	2.7236	0.3039	0.3473	0.3172	0.3610	0.2890	0.1103	0.2715	0.2564	0.3693	0.3785	0.8658
0.5	Average of Bias Error	-3.8002	-3.5901	-3.6008	-3.5322	-3.5202	-3.5957	-3.5348	-3.5475	-3.5679	-3.4869	-3.6864	-3.5869
	StdDev of Bias Error	2.3057	0.2388	0.4249	0.3360	0.4158	0.3152	0.1573	0.2771	0.2310	0.3529	0.2547	0.7392
0.6	Average of Bias Error	-4.5196	-4.6183	-4.2409	-4.3593	-4.1649	-4.2599	-4.2645	-4.2528	-4.3042	-4.1598	-4.2399	-4.3081
	StdDev of Bias Error	3.0054	0.3835	0.3560	0.4202	0.3714	0.2604	0.1841	0.3005	0.2719	0.3591	0.3567	0.9429
0.7	Average of Bias Error	-4.7062	-5.1416	-4.9396	-4.9997	-4.9693	-4.9323	-4.9486	-4.9444	-5.0147	-4.8859	-5.0030	-4.9529
	StdDev of Bias Error	2.9985	0.6291	0.4904	0.4123	0.4301	0.3267	0.1478	0.2903	0.2878	0.3834	0.2946	0.9552
0.8	Average of Bias Error	-6.3595	-5.9020	-5.7717	-5.7390	-5.6256	-5.6846	-5.6344	-5.6680	-5.6748	-5.6004	-5.6974	-5.7601
	StdDev of Bias Error	3.4464	0.5195	0.4698	0.4874	0.4426	0.3327	0.1771	0.3772	0.2864	0.3873	0.4642	1.0972
0.9	Average of Bias Error	-6.5465	-6.5061	-6.5109	-6.3733	-6.4132	-6.3845	-6.3011	-6.4033	-6.4456	-6.3190	-6.4665	-6.4243
	StdDev of Bias Error	3.4248	0.5222	0.4299	0.4181	0.5921	0.3609	0.2622	0.3001	0.2877	0.3768	0.4938	1.0776
1	Average of Bias Error	-7.3333	-7.1110	-7.1785	-7.2396	-7.0573	-6.9731	-7.0904	-7.1493	-7.1215	-7.0406	-7.1592	-7.1320
	StdDev of Bias Error	3.7353	0.5474	0.4847	0.5073	0.5070	0.3785	0.2914	0.3693	0.2931	0.4236	0.2948	1.1655
Total Average of Bias Error		-3.7263	-3.6467	-3.5920	-3.5787	-3.5120	-3.5415	-3.5527	-3.5518	-3.5766	-3.4871	-3.2708	-3.5481
Total StdDev of Bias Error		3.4502	2.3168	2.2980	2.3000	2.2850	2.2451	2.2204	2.2710	2.2678	2.2570	2.4125	2.4125

8. Bibliography

Allen, S., & D'Esopo, D. (1968). An Ordering Policy for Repairable Stock Items. *Institute for Operations Research and Management Science*, 16(3), 669–674.

Allen, S. G., & D'Esopo, D. A. (1968). An Ordering Policy for Repairable Stock Items. *Operations Research*, (3), 669.

AppleCare Protection Plan for Mac. (n.d.). Retrieved December 7, 2015, from <http://www.apple.com/support/products/mac.html>

Apple Vintage and Obsolete Products. (n.d.). Retrieved December 7, 2015, from

Bass, F. M. (1969). A New Product Growth for Consumer Durables. *Institute for Operations Research and Management Science*, 15(5), 215–227.

Berger, D. (2012). 5 Tips for Managing Spare Parts. Retrieved December 5, 2015, from <http://www.plantservices.com/articles/2012/02-asset-manager-managing-spare-parts/>

Blischke, W. R., & Murthy, D. N. P. (1992). Product warranty management - I: A taxonomy for warranty policies. *European Journal of Operational Research*, 62(2), 127–148.

Blumberg, D. F. (1999). Strategic examination of reverse logistics & repair service requirements, needs, market size, and opportunities. *Journal of Business Logistics*, 20(2), 141–159.

Brown, R. G. (1959). *Statistical Forecasting for Inventory Control*. McGraw-Hill.

Cohen, M. A., Agrawal, N., & Agrawal, V. (2006). Winning in the Aftermarket. *Harvard Business Review*, 84(5), 129–138.

-
- Cohen, M. A., & Agrawal, V. (1999). *After-Sales Service Supply Chains: A Benchmark Update of the North American Computer Industry* (p. 14). The Wharton School - University of Pennsylvania.
- Cohen, M. A., & Lee, H. A. (1990). Out of Touch With Customer Needs? Spare Parts and After Sales Service. *Sloan Management Review*, 31(2), 55.
- Croston, J. D. (1972). Forecasting and Stock Control for Intermittent Demands. *Operations Research Quarterly*, 23(3), 289–303.
- Dekker, R., Pince, C., Zuidwijk, R., & Jalil, M. N. (2013). On the use of installed base information for spare parts logistics: A review of ideas and industry practice. *International Journal of Production Economics*, 143, 536–545.
- Douthit, D., Flach, M., & Agarwal, V. (2011). A “Returning Problem”: Reducing the Quantity and Cost of Product Returns in Consumer Electronics. Retrieved November 1, 2015, from <http://www.cas-ag.org/SiteCollectionDocuments/PDF/Accenture-Reducing-the-Quantity-and-Cost-of-CustomerReturns.pdf>
- Douthit, D., Flach, M., & Agarwal, V. (2011). *Reducing the Quantity and Cost of Customer Returns*. Accenture Inc.
- Elwany, A. H., & Gebraeel, N. Z. (2008). Sensor-driven prognostic models for equipment replacement and spare parts inventory. *IIE Transactions*, 40(7), 629–639.
- Fleischmann, M., van Nunen, J. A., & Gräve, B. (2003). Integrating Closed-Loop Supply Chains and Spare-Parts Management at IBM. *Institute for Operations Research and Management Science*, 33(6), 44–56.
- Flowers, A. D., & O’Neill, J. B. (1978). An Application of Classical Inventory Analysis to a Spare Parts Inventory. *Interfaces*, 8(2), 76–79.

-
- Fortuin, L. (1980). The All-Time Requirement of Spare Parts for Service After Sales-Theoretical Analysis and Practical Results. *International Journal of Operations and Production Management*, 1(1), 59–70.
- Godoy, D., Pascual, R., & Knights, P. (2013). Critical spare parts ordering decisions using conditional reliability and stochastic lead time. *Reliability Engineering & System Safety*, 119(1), 199–206.
- Guide, D., Van Wassenhove, L. N., & Harrison, T. P. (2003). The Challenge of Closed-Loop Supply Chains. *Institute for Operations Research and Management Science*, 33(6), 3–6.
- Holt, C. C. (1959). Forecasting seasonals and trends by exponentially weighted moving averages. *International Journal of Forecasting*, 20(1), 5–10.
- Jalil, M. N., Zuidwijk, R. A., Fleischmann, M., & van Nunen, J. A. (2011). Spare parts logistics and installed base information. *The Journal of the Operational Research Society*, 62(3), 442–457.
- Jones, C. (2014). Apple's Warranty Accruals Mask True Gross Margins. Retrieved December 9, 2015, from <http://www.forbes.com/sites/chuckjones/2014/11/03/apples-warranty-accruals-mask-true-gross-margins/>
- Li, R., & Ryan, J. (2011). A Bayesian Inventory Model Using Real-Time Condition Monitoring Information. *Production and Operations Management*, 20(5), 754–771.
- Louit, D., Pascual, R., Banjevic, D., & Jardine, A. K. S. (2011). Condition-based spares ordering for critical components. *Mechanical Systems and Signal Processing*, 25, 1837–1848.

-
- McGuire, E. P. (1980). *Industrial Product Warranties: Policies and Practices*. New York: Conference Board.
- Molenaers, A., Baets, H., Pintelon, L., & Waeyenbergh, G. (2012). Criticality Classification of Spare Parts: A Case Study. *International Journal of Production Economics*, 140(2), 570–578.
- Moore, J. R. (1971). Forecasting and Scheduling for Past-Model Replacement Parts. *Institute for Operations Research and Management Science*, 18(4), B200–B213.
- Murthy, D. N. P., Solem, O., & Roren, T. (2004). Product warranty logistics: Issues and challenges. *European Journal of Operational Research*, 156(1), 110–126.
- Overton, D. (2005). “No Fault Found” returns cost the mobile industry \$4.5 billion per year. Retrieved December 12, 2015, from <http://www.wds.co/no-fault-found-returns-cost-the-mobile-industry-4-5-billion-per-year/>
- Porter, M. E., & Heppelmann, J. E. (2015). How Smart, Connected Products are Transforming Companies. *Harvard Business Review*, 93(10), 96–114.
- Priest, G. L. (1981). A Theory of the Consumer Product Warranty. *The Yale Law Review*, 90(6), 1297–1352.
- Qi, H., Ganesan, S., & Pecht, M. (2008). No-fault-found and intermittent failures in electronic products. *Microelectronics Reliability*, 48(1), 663–674.
- Ritchie, E., & Wilcox, P. (1977). Renewal theory forecasting for stock control. *European Journal of Operational Research*, 1(2), 90–93.
- Schrady, D. (1967). A Deterministic Inventory Model for Repairable Items. *Naval Research Logistics Quarterly*, 14, 391–398.

-
- Solomon, R., Sandborn, P., & Pecht, M. (2000). Electronic Part Life Cycle Concept and Obsolescence Forecasting. *IEEE Transactions on Components and Packaging Technologies*, 23(4), 707–718.
- Syntetos, A. A., & Boylan, J. E. (2001). On the bias of intermittent demand estimates. *International Journal of Production Economics*, 71(1-3), 457–466.
- Teunter, R. H., & Fortuin, L. (1999). End-of-life service. *International Journal of Production Economics*, 59(1), 487–497.
- Teunter, R. H., & Haneveld, W. K. K. (2002). Inventory control of service parts in the final phase. *European Journal of Operational Research*, 137(1), 497–511.
- Toshiba Extended Service Plans. (n.d.). Retrieved December 7, 2015, from <http://us.toshiba.com/serviceplans>
- Tracht, K., Goch, G., Schuh, P., Sorg, M., & Westerkamp, J. (2013). Failure probability prediction based on condition monitoring data of wind energy systems for spare parts supply. *CIRP Annals - Manufacturing Technology*, 62(1), 127–130.
- Willemain, T. R., Smart, C. N., & Schwarz, H. F. (2004). A new approach to forecasting intermittent demand for service parts inventories, 20(3), 375–387.
- Winter, P. R. (1960). Forecasting Sales by Exponentially Weighted Moving Averages, 6(3), 324–342.