

A BAYESIAN APPROACH TO DEMAND FORECASTING

A Thesis Presented to the Faculty of the Graduate School

University of Missouri-Columbia

In Partial Fulfillment

Of the Requirements for the Degree

Masters of Science

By

JENNIFER JEAN BERGMAN

Dr. James S. Noble, Thesis Advisor

Dr. Ronald McGarvey, Thesis Co-Advisor

December 2014

The undersigned, appointed by the Dean of the Graduate Faculty, have examined the thesis entitled

A BAYESIAN APPROACH TO DEMAND FORECASTING

Presented by Jennifer Jean Bergman,

A candidate for the degree of Master of Science

And hereby certify that, in their opinion, it is worthy of acceptance.

Dr. James S. Noble, Thesis Advisor

Dr. Ronald McGarvey, Thesis Co-Advisor

Dr. Timothy Matisziw

DEDICATION

I would like to dedicate this research to my dad, Larry, and my mom, Judy. My dad has been such an amazing role model and has always supported me in every decision I make. He is extremely selfless and has taught me how to be a great yet kind leader. My mom has gone out of her way to include me in all of her business ventures so that I can learn how to be successful early in life, and every action she makes is always in the best interest of me. I am so blessed to have parents like you and always want to make you guys proud. Thank you for being such a big part of my achievements.

ACKNOWLEDGEMENTS

I owe the utmost gratitude to all the people who have supported me throughout the process of my thesis. My deepest gratitude goes to my advisor, Dr. James Noble. He has always believed in me, and without him I would not have considered a graduate career in Industrial Engineering. Dr. Noble gave me the opportunity to join his research team, and from this positive experience, I decided to pursue a Master's degree. He has truly guided and supported me throughout my entire college career, and I really appreciate all of the time he has put into the writing process of my thesis. I would like to thank my co-advisor, Dr. Ron McGarvey. He is always so willing to help me find solutions to all of my difficult research questions, and I am very appreciative of his encouragement and direction throughout my thesis project. Also, I must comment on how much I admire his attention to detail, which has aided in successful research presentations. Without Dr. Noble or Dr. McGarvey, my research would not have been possible.

A very special thanks goes to my entire Center of Excellence in Logistics and Distribution (CELDi) research team. Foremost, I would like to acknowledge Randolph Bradley for all of his hard work and time towards helping with my research. I am truly grateful for his interest and assistance in making my thesis project successful, and I have enjoyed all of our long phone meetings. Randolph is my first professional mentor, and his excitement in all areas of work and life is inspiring. I would also like to extend a thank you to Steve Saylor for all of his help on the simulation model used in this thesis. I had many model modification requests, and he always responded in an extremely timely manner.

Last, I must acknowledge all of my friends and family. All of their love and support has really contributed to my success in graduate school, and I am so fortunate to have such positive and motivating people in my life.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	ii
LIST OF FIGURES	vi
LIST OF TABLES	vii
ABSTRACT	viii
Chapter 1 – Introduction.....	1
1.1 Overview of Demand Forecasting within Supply Chain Management	1
1.2 Motivation and Current Practices.....	2
1.3 Introduction to Applying Bayesian Estimation to Demand Forecasting	6
1.4 Summary of Research Objectives	7
Chapter 2 – Literature Review.....	9
2.1 Overview of Relevant Literature	9
2.2 Judgmental Forecasting	10
2.3 Statistical Forecasting	11
2.3.1 Time Series Forecasting Models	12
2.3.2 Causal Forecasting Models	14
2.3.3 Bootstrap Forecasting Models	15
2.3.4 Neural Network Forecasting Models	15
2.3.5 Statistical Forecasting Conclusion.....	16
2.4 Bayesian Forecasting.....	16
2.4.1 Bayesian Background	17
2.4.2 Review of Bayesian Literature	22
2.5 Critique of Literature	27
Chapter 3 – Problem Description and Model Formulation	34
3.1 Problem Overview	34
3.2 Current Demand Forecasting Practices	34
3.3 Formulating Demand Forecasting Using Bayes’ Rule Equation	38
3.3.1 Model Assumptions	38
3.3.2 Formulating Bayesian Model	40
3.4 Problem Description and Model Formulation Summary	45
Chapter 4 – Case Study Methodology	47

4.1 Overview of Case Study Methodology.....	47
4.2 Overview of Networks	47
4.3 Impact on Forecast Accuracy	49
4.3.1 Forecast Accuracy Metrics	49
4.3.2 Impact on Forecast Accuracy Methodology	51
4.4 Impact on Inventory.....	53
4.4.1 Strategic Inventory Optimization	53
4.4.2 Impact on Inventory Methodology	55
4.5 Impact on Overall Supply Chain Performance	56
4.5.1 Discrete-Event Warehouse Simulation	56
4.5.2 Model Validation.....	57
4.5.3 Impact on Supply Chain Performance Methodology.....	59
Chapter 5 – Case Study Results	64
5.1 Overview of Case Study	64
5.2 Impact on Forecast Accuracy	64
5.2.1 Forecast Accuracy	65
5.2.2 Low versus High Demand Parts	67
5.2.3 Appropriate Weighting Scheme.....	70
5.2.4 LRU Forecast Accuracy	71
5.2.5 Forecast Accuracy Conclusion.....	72
5.3 Impact on Inventory.....	73
5.3.1 Mean Time between Demands versus Stock Levels	73
5.3.2 Inventory Investment Costs	76
5.4 Impact on Fill Rate Performance	78
5.4.1 Model Validation.....	78
5.4.2 Fill Rate Performance versus Inventory Investment.....	80
5.4.3 Fill Rate versus Forecast Accuracy	83
Chapter 6 – Conclusions, Contributions and Directions for Future Research	85
6.1 Conclusions / Contributions.....	85
6.2 Directions for Future Research	88
6.2.1 Hybrid Method.....	88

6.2.2 Analysis on Different Scenario	88
6.2.3 Current Data Limitations.....	89
REFERENCES.....	90

LIST OF FIGURES

Figure 2-1 Confidence Interval versus Credible Interval (Recchia, 2012)	19
Figure 4-1 Part Breakdown by Network	48
Figure 4-2 Current Inventory Investment by Network	49
Figure 4-3 Simulation Model Validation Methodology	57
Figure 4-4 Cumulative Histogram for Demands per Month	61
Figure 5-1 Optimal Inventory Investment Costs	76
Figure 5-2 Chart of fill rate over time, with an 80% confidence interval, for Bayesian Method (Green), Current Method (Orange) and Engineering Estimates (Red). Higher fill rate is better.	81
Figure 5-3 Fill Rate Performance Comparison of Current Approach versus Bayesian Approach	81

LIST OF TABLES

Table 3.1 Level of Confidence in Engineering Estimate for Different Parts Groups	35
Table 3.2 Blended Forecast Equation Notation	37
Table 3.3 Part Demand Dates	38
Table 3.4 Aggregated Flight Hours per Month	39
Table 5.1 Forecast Accuracy in Year 1	65
Table 5.2 Forecast Accuracy in Year 2	65
Table 5.3 Forecast Accuracy in Year 3	65
Table 5.4 Forecast Accuracy in Year 4	66
Table 5.5 Forecast Accuracy if Outlier Part was Weighted Appropriately in Year 2.....	67
Table 5.6 Forecast Accuracy for Low Demand versus High Demand Parts	68
Table 5.7 Forecast Accuracy for All Low Demand Parts.....	68
Table 5.8 Forecast Accuracy for Low Demand Parts with Demand.....	69
Table 5.9 Forecast Accuracy for all High Demand Parts	69
Table 5.10 Forecast Accuracy for High Demand Parts with Demand	69
Table 5.11 Forecast Accuracy Summary for Source of Engineering Estimate	70
Table 5.12 Forecast Accuracy Detailed for Source of Engineering Estimate	71
Table 5.13 Summary of Forecast Accuracy of LRU Parts.....	72
Table 5.14 Forecast Accuracy of LRU Parts.....	72
Table 5.15 Mean Time between Demands (MTDB) versus Target Stock Levels (TSL)	74
Table 5.16 Bayesian Method versus Engineering Estimate Mean Time between Demands (MTBD) and Target Stock Levels (TSLs).....	74
Table 5.17 Current Method versus Engineering Estimate Mean Time between Demands (MTBD) and Target Stock Levels (TSLs).....	75
Table 5.18 Bayesian Method versus Current Method Mean Time between Demands (MTBD) and Target Stock Levels (TSLs)	75
Table 5.19 Simulated Cumulative Fill Rate over Ten years.....	79
Table 5.20 Student's t-Test Inputs for the Null Hypothesis that Simulated Fill Rate is Equal to the Optimization Fill Rate of 80.58%.....	80
Table 5.21 Student's t-Test Outputs for the Null Hypothesis that Simulated Fill Rate is Equal to the Optimization Fill Rate of 80.58%.....	80
Table 5.22 Fill Rate versus Forecast Accuracy	84

A BAYESIAN APPROACH TO DEMAND FORECASTING

Jennifer Jean Bergman

Dr. James Noble, Thesis Advisor

Dr. Ronald McGarvey, Thesis Co-Advisor

ABSTRACT

Demand forecasting is a fundamental aspect of inventory management. Forecasts are crucial in determining inventory stock levels, and accurately estimating future demand for spare parts has been an ongoing challenge, especially in the aerospace industry. If spare parts are not readily available, aircraft availability can be compromised leading to excessive downtime costs. As a result, inventory investment for spare parts can be significant to ensure down time is minimized. Additionally, most aircraft spare parts are considered 'slow-moving' and experience intermittent demand making the use of traditional forecasting methods difficult in this industry. In this research, a forecasting method is developed using Bayes' rule to improve the demand forecasting of spare parts. The proposed Bayesian method is especially targeted to support new aircraft programs and is not intended to change how inventory is currently optimized. A case study based on a real aircraft program's data is performed in order to validate the use of the proposed Bayesian method. In the case study, three forecasting methods are compared: judgmental forecasting, a traditional statistical forecasting approach, and the proposed Bayesian method. The methods' impact on forecast accuracy, inventory costs, and fill rate performance (evaluated using simulation) are analyzed. The results conclude that the proposed Bayesian approach outperforms the other methods in terms of fill rate performance. Hence, the Bayesian method improves demand prediction and thus, more accurately estimates inventory needs allowing managers to make better inventory investment decisions.

Chapter 1 – Introduction

1.1 Overview of Demand Forecasting within Supply Chain Management

Due to today's competitive and global economy, there has been a substantial shift on how companies view their supply chains. In recent years, companies have invested a significant amount in their supply chain management programs in order to attain a competitive edge and protect their market share (Li et al., 2006). There are numerous key functions that can impact supply chain management, which have received ample interest in both academia and industry. Some of these key functions that have received attention in recent research are inventory management, transportation management, sourcing and marketing in a supply chain, green logistics, and behavior operations (Li, 2014). This thesis will focus primarily on inventory management.

Until recently the benefits of effective inventory management were not well understood. The questions answered in inventory management seemed straightforward and overstocking seemed advantageous especially in terms of reducing stock outs. Conversely, in the context of reality, inventory concerns are ambiguous, and balancing inventory is extremely important. Inventory accounts for approximately one-third of all assets in a company (Loui et al., 2011) and financing this inventory is both challenging and expensive. The trade-off is clear: a small amount of inventory can result in poor supply support or costly penalties, and a large amount of inventory can result in non-value added capital tied up in unused inventory. These inventory stock levels must also stay within budget while meeting customer performance metrics. For these reasons, it is apparent that inventory management is challenging and important to strategically plan.

Moreover, inventory investment decisions are usually determined months before demand is needed. This means that managers are making extremely important and complex investment decisions in the face of uncertainty. For these reasons, it is important to effectively manage uncertainty in order to improve inventory control. However, the big question is, “how can a company cope with this natural phenomena of uncertainty?” It is evident that improving demand forecasts sequentially minimizes uncertainty and thus, improves receptivity to volatile demand. Increasing responsiveness to fluctuating demand balances inventory and in turn, streamlines a supply chain. The accuracy of a company’s forecast has been proven to have a significant impact on the performance of a supply chain (Ton de et al., 2005). From this it is clear that demand forecasting drives inventory control and is an extremely important player in supply chain management.

1.2 Motivation and Current Practices

Demand forecasting is an essential component in inventory management. Inventory stock levels are dependent on forecasts of demand, and accurately estimating spare part’s demand in the aerospace industry has been a continuing concern. Aircraft unavailability can lead to significant downtime costs, and the lead-time to repair certain parts can be extensive. As a result, inventory investment for spare parts can be excessive to minimize downtime. Additionally, most spare parts are considered ‘slow-moving’ and experience intermittent demand. Intermittent demand occurs at infrequent times with long periods without demand, which creates difficulties when using traditional statistical methods to forecast demand. In order to illustrate the potential opportunities in spare parts management, Doug Blazer (1996) summarized that in 1996 the United States Air Force maintained over 30 billion dollars in reparable spare parts inventory and spent between 2 to 3 billion a year to repair them (Future Vision for the Air Force Logistics System – Doug Blazer). From this it is clear that this has been a

continuing concern in the aviation industry, and there is opportunity for improvement in this field. In order to further motivate this research, current practices to forecast demand are first reviewed. Second, the Performance-Based Logistics contract is explained, and this contract emphasizes the need to more accurately forecast demand. Last, research addressing spare parts demand forecasting is briefly summarized and used to motivate certain objectives in this thesis.

First, consider the current practices for forecasting demand for spare parts on a new aircraft program. The most common approach to forecast demand is to utilize statistical methods such as simple exponential smoothing. However, these approaches require observed demand data, and when a new program is employed, no such historical information exists requiring the use of engineering estimates. These estimates approximate the mean time between failures and can come from a variety of sources. For example, some estimates are based on large amounts of historical demand experience on similar aircraft programs, which increase manager's confidence in these estimates. Managers are less confident in estimates for parts unique to a particular aircraft, and that are based on an engineer's best guess. However, no matter how confident one is in the estimate, all initial estimates of demand contain a considerable amount of uncertainty, and programs rely heavily on these engineering estimates to evaluate optimal stock levels. Also, companies struggle with systematically incorporating one's confidence in these estimates when forecasting demand. For these reasons, this thesis explores the risk associated with the uncertainty of engineering estimates and seeks to create a methodology that accounts for one's confidence in engineering estimates.

Next, after employing a new aircraft program, the big question is when the program should transition from engineering estimates to statistical approaches using observed demand.

There are currently a couple of common techniques used to make this transition. First, this transition can occur once a sufficient amount of operating hours have developed. Next, this transition can occur once a part experiences a specified amount of accumulated demand. Both of these transitions can be problematic. For example, suppose the transition from engineering estimates to statistical methods using actual demand occurs once 100,000 operating hours have occurred. On certain aircraft programs, this amount of operating hours can take at least five years to develop, which is a long period of time to optimize inventory based on estimates that could be poor forecasts of demand. Thus, it is clear that there is opportunity to explore other methods that evaluate or eliminate the transition from engineering estimates to observed demand. Additionally, it is important to note that once transitioning to actual demand, a blend of simple exponential smoothing and causal factors is used to estimate demand. However, this statistical approach experience difficulties when predicting demand for intermittent and 'slow-moving' parts. This will be discussed further in Chapter 2.

Second, the concept of a Performance-Based Logistics (PBL) contract is explained. Aircraft manufacturers are contracted to provide considerable support for new aircraft programs through a PBL contract. A report by the Center for the Management of Science and Technology at the University of Alabama in Huntsville defines PBL as,

“An integrated acquisition and sustainment strategy for enhancing weapon system capability and readiness, where the contractual mechanisms will include long-term relationships and appropriately structured incentives with service providers, both organic and non-organic, to support the end user’s (warfighter’s) objectives” (Berkowitz et al., 2003).

In other words, aircraft manufacturers are contracted to supply spare parts in order to achieve an agreed upon fill rate or aircraft availability goal. In order to supply spare parts, aircraft manufacturers must determine the right mix of parts to stock in order to achieve the specified operational metric. This mix of parts is a function of predicted demand and an agreed upon

service level. However, as expressed above, forecasting demand for spare parts in this industry can be extremely difficult due to low and intermittent demand. From this it is clear that aircraft manufacturers must account for this issue of low and intermittent demand while achieving a specified operational metric and minimizing inventory costs.

Last, it is important to briefly note the current literature on demand forecasting for spare parts in order to motivate a major objective in this thesis. Literature in this field has recently gained increased attention due to researcher's interest in lumpy or intermittent demand (Bacchetti and Sacconi, 2012). This indicates that new research in this area is very attractive. Additionally, a gap exists between research and practice in the field of spare parts management (Bacchetti and Sacconi, 2012; Cohen et al., 2006; Wagner and Lindemann, 2008). The reasons behind this will be discussed further in Chapter 2. However, due to the gap between research and practice, this research focuses on developing a methodology that can be easily implemented in industry, which is an extremely important objective throughout this research. Additionally, a case study applied to a real aircraft program will be performed after developing this methodology. The case study will benchmark how an actual aircraft manufacturer is currently forecasting demand of spare parts and will validate the use of the proposed approach in real application. Overall, this thesis aims to help close the gap between research and practice by providing a practical demand forecasting solution and performing a real industrial case study.

In conclusion, it is evident that there is substantial motivation to develop a new method to forecast demand of spare parts, especially for new aircraft programs. First, demand for aircraft spare parts is low and intermittent making the use of traditional forecasting methods difficult in this industry. Hence, program managers rely heavily on engineering estimates. However, there is a lot of uncertainty when using engineering estimates, so it is important to illustrate these

risks and propose a model that does include initial confidence in these estimates. Second, there is support to establish a methodology that improves or eliminates when parts will transition from engineering estimates to observed demand. Third, aircraft manufacturers typically support new aircraft programs through a PBL contract. This means that they must determine the right mix of parts to stock in order to achieve the specified operational metric and minimize costs making the problem even more complex. Last, literature on demand forecasting for spare parts is young, and there is a gap between research and practice in the field of spare parts management. This motivates continued research in this field and the development of a specific demand forecasting recipe that can be used in practice.

1.3 Introduction to Applying Bayesian Estimation to Demand Forecasting

Demand estimates for spare parts are currently based on engineering estimates or observed demand. The majority of these replacement parts incur little to no demand, so an approach that utilizes all available information could be very advantageous, especially for new equipment programs. The Bayesian methodology provides an intelligent way of combining prior knowledge with observed data to obtain a new and improved estimate of demand. This framework has been justified for estimating demand for parts that lack failure observations (Sherbrooke, 2004), which has been a significant motive for the use of Bayesian estimation in industrial application. Additionally, Bayesian estimation has the ability to characterize one's uncertainty through probability statements, which has been another major reason for using Bayes' in practice. Thinking in terms of probability is intuitive and can be an extremely useful tool in decision-making. This decision-making feature has made Bayesian analysis extremely popular in a variety of fields from medical diagnosis to machine learning. For these reasons, Bayes' rule for demand forecasting is extremely worth investigating.

It is important to note that Bayes' theorem has a long history with much controversy dating back to the 1700s. There are two major paradigms in mathematical statistics: frequentist and Bayesian. The basis of controversy is how the Bayesian approach views the nature of probability and parameters compared to frequentist statistics. However, Bayesian estimation has proven effective in many practical applications making this approach appealing. The differences between Bayes' and frequentist statistics will be explained in Chapter 2.

1.4 Summary of Research Objectives

Demand estimation has a significant impact on inventory control and supply chain management. Hence, improving estimates of spare parts demand in the aerospace industry has been an area of increasing interest. Most spare parts are considered 'slow-moving' and experience intermittent demand. This makes predicting demand for spare parts much more challenging. Thus, program managers rely heavily on engineering estimates. These estimates are used to forecast demand during the early life of a program until enough observed data is accumulated. However, if these estimates are poor estimates of demand and do not accurately depict how a part is truly behaving, poor performance will result. Also, these estimates come from a variety of sources, and there is currently no method to include one's initial confidence in these estimates. Thus, a method that can utilize these initial estimates of demand while learning from observed data could significantly improve performance, especially during the early life of a program.

This research utilizes the Bayesian forecasting approach to update prior information as observed demand occurs. The prior information makes up for initial variation in data while continuously learning from actual demand. This research explores the use of Bayes' rule to improve the demand forecasting of spare parts for new equipment programs. It is important to

note that the proposed approach is not intended to change how inventory is optimized. It is to improve demand prediction in order to more accurately optimize stock levels and allow managers to make better inventory investment decisions.

The remainder of this thesis is organized as follows. Chapter 2 describes existing literature on how others have forecasted demand for spare parts and utilized Bayes' rule to forecast demand. This discussion emphasizes the importance of continued research in this area and prospects for addressing current shortcomings. Chapter 3 further details the problem described in the introduction and formulates the demand forecasting equation using Bayes' rule. After applying Bayes' rule to demand forecasting, a case study on a real aircraft program's data is presented to validate the use of this new method. Chapter 4 presents the methodology of the case study, and Chapter 5 provides the results of the case study. Last, Chapter 6 states the conclusions from this research and recommends directions for future research.

Chapter 2 – Literature Review

2.1 Overview of Relevant Literature

Literature on operations management in the face of demand uncertainty dates back to the 1950s (Arrow, 1958), and this awareness paved the evolutionary path for demand forecasting. Research in demand forecasting is well developed and continues to be an interest in both academia and industry. More specifically, demand forecasting of spare parts has recently gained increased attention. Over the past few decades, maintenance has become increasingly important in industrial environments, and this area of research is growing at an extraordinary rate (Callegaro, 2009). Effective maintenance is dependent on spare parts availability, and due to this relationship, research in spare parts management is advancing as well. Additionally, ample concern in intermittent and lumpy demand has expanded literature in the specific area of spare parts forecasting.

Many reviews have expressed a research-practice gap in the study of spare parts management. Adrodegari et al. (2014) is currently performing a critical review on spare parts inventory management. They conclude in their preliminary observations that there are a limited number of papers that give a practitioner's view on how to apply proposed methods. Only 24% of the 191 papers relating to this topic include empirical applications through case studies. This could explain the research-practice gap discussed in previous literature (Boone et al., 2008; Syntetos, Keyes, et al., 2009; Wagner and Lindemann, 2008). Additionally, forecasting methods in research can seem too complex for practitioners, too costly to integrate with their current systems, or include assumptions that are not realistic for the real world. This illustrates a need to bridge the gap between research and practice.

Overall, it is clear that demand forecasting has been a continuing area of interest. Research in forecasting methods to estimate spare part demand has gained renewed attention and is excelling due to concern in maintenance modeling and intermittent demand. Although the theoretical knowledge has sufficiently increased, the literature has identified a gap between research and practice. This section will review three types of forecasting methods researched and utilized in spare parts management: judgmental forecasting, statistical forecasting, and Bayesian forecasting.

2.2 Judgmental Forecasting

Judgmental forecasts are formed by expert opinion and very common in practice (Klassen and Flores, 2001; McCarthy et al., 2006). When little to no historical data is available, judgmental methods may initially be used and can often exhibit positive results. This type of forecasting is also frequently used as an adjustment method. For example, managers are found to use their opinion as a correction factor to statistical techniques. Such modifications can provide better accuracy, and there has been some empirical research on the merits of using judgmental forecasting along with statistical methods. The remainder of this section will briefly summarize why judgmental forecasting is common in practice and discuss some empirical analysis when combining expert opinion with statistical analysis.

Goodwin (2002) points out many reasons for the prevalence of judgmental forecasting in industry. First, this type of forecasting is utilized when there is a limited amount of historical demand as required by many statistical methods. Second, it can be used when statistical models cannot exhibit effects of special events that may influence the future. Third, it is often used when modelers have a lack of understanding behind the statistical methods (also known as the “black box” effect). Although this type of forecasting is common in practice, there is a

subjective bias when utilizing opinion, and it is important to be cognizant of this in order to minimize subjective error. Goodwin (2002) concludes that frequency of unnecessary judgmental adjustments can be lowered when forecasters provide rationale for making these adjustments.

Most of the support for judgmental adjustments is in economic forecasting literature (i.e. Turner, 1990). Only a few studies have explored the advantages of judgmental adjustments in the context of demand forecasts for SKUs. For example, Syntetos, Nikolopoulos, et al. (2009) analyzes monthly intermittent demand forecasts for a major international pharmaceutical company in the United Kingdom. In their case study, the company uses a statistical forecasting system to forecast demands that are successively adjusted by judgment based on marketing data. They conclude that there is benefit from judgmental adjustments when parts are slow-moving and intermittent. This was the first study to publish evidence on the effectiveness of opinion in forecasts for intermittent demand. However, this effectiveness is dependent on the nature of the adjustments and characteristics of the time series.

Overall, judgmental forecasting is extremely common in practice. There are many reasons for this summarized as summarized by Goodwin (2002), and this forecasting approach can yield beneficial in some instances. Nevertheless, further research is needed on the implications of forecasting adjustments on prediction and inventory (Boylan and Syntetos, 2009).

2.3 Statistical Forecasting

Statistical forecasting can be used when historical data is available. There are numerous statistical forecasting methods discussed in literature. However, this section focuses primarily

on the literature related to spare parts forecasting. Four statistical forecasting approaches are discussed: time-series, causal factors, bootstrapping, and neural networks.

2.3.1 Time Series Forecasting Models

A time series is a collection of historical data in a given period of time. Time series forecasting methods find patterns in data to predict the future. Traditional time series methods such as moving average and simple exponential smoothing (also known as single exponential smoothing or SES for short) are frequently used in practice. Moving average assumes all periods have equal importance. However, when all past observations should not be weighted the same, weighted moving average (WMA) where weights can be assigned to the most recent observations or simple exponential smoothing where the value of an observation degrades over time can be used. Brown (1959) invented SES while he was an operations research analyst for the US Navy. His methodology revolves around practicality constraints and is easy to implement making it remarkably attractive in practice. SES is also largely used to forecast in inventory control systems (Syntetos, Boylan, et al., 2009). However, there are some disadvantages when using this approach. The smoothing constant can be difficult to choose. If the constant is small, then response to change is slow. If the constant is large, then the response to change is fast; however, the output can include a large amount of variability. Also, along with moving average methods, if a trend in the data exists, these methods lag. For example, if the mean is steadily increasing, then the forecast will be several periods behind.

Holt (1957) refines Brown's SES model by adding a trend smoothing constant to account for trends in data. Many have illustrated that Holt's method (known as double exponential smoothing) works well with problems that incorporate trend. However, Brown's approach is still recommended in practice if trend and seasonality do not exist. It is important to add that

Holt's approach does not include seasonality. Consequently, one of Holt's students, Winters (1960), extended Holt's method to include a form of exponential smoothing that incorporates both trend and seasonality. This approach is known as the Holt-Winter method or triple exponential smoothing. Although these approaches can handle seasonality and trend issues, these methods add additional smoothing constants that could increase difficulty when initiating forecasts. The choices for the smoothing constants are critical, and an inadequate choice could result in erroneous results. Additionally, the exponential smoothing methods discussed do not recognize that there can be periods in which zero demand for a part exists, which is very common for spare parts.

In order to cope with this issue of periods with zero demand, Croston (1972) proposed a simple exponential smoothing technique that updates forecasts only in periods with demand. His method first separates the size of non-zero demands and the interval between these demands when exponentially smoothing. Next, demand per period is estimated based upon the ratio of these estimates (size/interval). Most research in the field of intermittent demand forecasting is based on the seminal work done by Croston and his approach has been implemented in leading software packages for statistical forecasting (Boylan and Syntetos, 2009).

Croston's work was theoretically superior; however, some empirical evidence demonstrates the method's results as modest (Willemain et al., 1994) or worse than simpler methods (Sani and Kingsman, 1997). Due to this, there are many variants of Croston's work (Johnston and Boylan, 1996; Levén and Segerstedt, 2004; Syntetos and Boylan, 2005). An influential work done by Syntetos and Boylan (2001) illustrates if inventory rules consider expected demand per period together, then a bias will result when using Croston's method.

Later, Syntetos and Boylan (2005) modify Croston's approach by introducing an unbiased estimator. This method is known as Syntetos-Boylan Approximation (SBA). They performed a simulation experiment to compare the methods and conclude that Croston's method is only superior when the smoothing constant is low (less than .15). Other empirical studies following their work indicate that the SBA estimator reduces stock-holding costs while attaining a specified service level (Eaves and Kingsman, 2004; Gutierrez et al., 2008).

Ghobbar and Friend (2003) perform a comparative study between thirteen forecasting methods (most being time-series) to predict spare part demands for airline fleets. Some of the compared methods include weighted moving average, single exponential smoothing, Holt's method, Holt-Winter's method, and Croston's method. They state that SES and Mean Time between Replacement (MTBR) are the most common forecasting methods used for airline operators. In the comparative study, analysis of variance (ANOVA) is used to evaluate variation, and Mean Absolute Percentage Error (MAPE) is used to evaluate forecast accuracy. The experimental results show that weighted moving average performs superior to the other models, and Holt-Winter's method performs the worst. Overall, from this study they determine that the commonly used forecasting methods (SES and MTBR) in the airline industry can be questionable compared to other methods like weighted moving average.

2.3.2 Causal Forecasting Models

Boylan and Syntetos (2008) explain that causal methods can be used during the initial life of a part when a reasonable amount of historical data is not yet available for statistical methods. Causal models assume that demand (dependent variable) has a cause-and-effect relationship with one or more independent variables, and the independent variable(s) must be specified before predicting demand. This method is commonly utilized in practice when the

aircraft fleet is new and increasing. Causal models can easily compute estimates of demand (dependent variable) with increased expected flight hours (independent variable). However, one criticism of this approach is that it assumes that the dependent variable is influenced solely by the independent variable(s) specified, which might not be the case.

2.3.3 Bootstrap Forecasting Models

Hill et al. (1996) points out that the traditional time-series models can misjudge the functional relationship between independent and dependent variables. In order to address this, a bootstrapping method is proposed. This approach has received much interest (and criticism) in academia and is commonly researched for forecasting intermittent items (Syntetos, Boylan, et al., 2009). Willemain et al. (2004) developed a patented heuristic to forecast intermittent demand for service parts using a bootstrapping approach. They were the first to suggest a modified bootstrap technique that forecasts the cumulative distribution of demand over a fixed lead-time. Their method combines bootstrapping, autocorrelation, and jittering. They compare the forecast accuracy of their approach against simple exponential smoothing and Croston's method on nine industrial datasets. Their results indicate that the bootstrapping method can lead to superior results. There are many variants of bootstrapping approaches (Bookbinder and Lordahl, 1989; Efron, 1979; Porras Musalem, 2005); however, there is no conclusive evidence that bootstrapping outperforms general methods.

2.3.4 Neural Network Forecasting Models

Neural networks are able to capture linear or non-linear relationships, which traditional time-series methods have difficulties in doing. Gutierrez et al. (2008) created a neural network model to forecast lumpy demand. They utilize neural networks to approximate the functional relationships within the data. Three forecast accuracy metrics are used to compare their model

to three time-series models (SES, Croston's method, and SBA). The results show that their approach performs significantly better than the traditional time-series methods. They also validate the results from Syntetos and Boylan (2005) that SBA outperforms SES and Croston's method when demand is lumpy. Additionally, Li and Kuo (2008) express that traditional neural networks have the opportunity for increased forecasting accuracy, so they use an enhanced fuzzy neural network (EFNN) for forecasting the demand for spare parts in the automobile industry. They conclude that the EFNN model outperforms the traditional neural networks in both fill rate and stock costs. Although neural networks perform competitively, there are still many practical limitations. A large amount of training data is needed, and the proposed neural networks are specific to the scenario making it hard to generalize these models for other applications.

2.3.5 Statistical Forecasting Conclusion

There are advantages and disadvantages of all of the statistical forecasting approaches discussed, and there is no conclusive results determining what method is best. An important underlying assumption in all of the models is that past behavior represents the future. However, when starting a new aircraft program, initial variation can cause early forecasts to misjudge the future. It could be advantageous to use observed information with expert opinion to forecast demand especially when data is limited.

2.4 Bayesian Forecasting

Judgmental input can prove especially valuable for slow-moving parts, and observed demand can validate how these parts are truly behaving. The Bayesian approach is a methodology that utilizes all available information. It is a systematic approach that updates prior information (such as judgmental input) as observed demand occurs. The Bayesian

paradigm is commonly used to overcome difficulties with limited demand data (Graves et al., 1993; Sherbrooke, 2004), and limited data is very common for aircraft spare parts. When new programs are employed, there is little to no historical data. Nevertheless, even when programs have been operating for several years, numerous parts are designed for reliability and will take years before demand is needed. The obstacles associated with limited demand data when forecasting the future need for spare parts makes Bayesian forecasting worth investigating. This section will delve into Bayesian forecasting by first providing background on Bayesian Statistics in the context of this research problem and then reviewing existing research on demand forecasting using Bayes' rule

2.4.1 Bayesian Background

Mathematical statistics is based on two major paradigms: frequentist and Bayesian. Bayesian statistics has a long history with much controversy dating back to the 1700s. An exemplary history of Bayes' rule and example of how this rule was utilized in searching for a lost submarine can be found in Sharon McGrayne's book, *The Theory That Would Not Die* (McGrayne, 2011). However, due to the length of history, this sub-section will begin with a brief summarization on the differences between the Bayesian approach and frequentist approach, which has been the basis of controversy throughout history. Then it will explain Bayes' rule and methods to solve this rule. This will allow the reader to have a full understanding of Bayesian statistics before reviewing previous literature.

To begin, the key differences between Bayesian and frequentist statistics will be discussed. The most obvious difference between the two is that Bayesian statistics includes prior knowledge and observed data where frequentist statistics only utilizes observed data. Implementing prior knowledge can be subjective, making it an area of criticism in statistics.

Next, there are differences in the nature of probability, parameters, and inferences. O’Hagan and Luce (2003) summarize these key differences shown in Table 2.1. Additionally, to expand on the differences in inference, frequentist statistics uses a confidence interval where Bayesian statistics uses a credible interval. A confidence interval is like a frequency statement. It takes random samples from a population and concludes that some percent of the samples contain the true parameter. A credible interval is a probability statement. It concludes that there is some percent chance that the interval contains the true parameter. Figure 2.1 is a great illustration on the differences between the two intervals (Recchia, 2012).

Table 2.1– Summary of Key Differences between Frequentist and Bayesian Statistics (O’Hagan and Luce, 2003)

FREQUENTIST	BAYESIAN
<i>Nature of probability</i>	
Probability is a limiting, long-run frequency.	Probability measures a personal degree of belief.
It only applies to events that are (at least in principle) repeatable.	It applies to any event or proposition about which we are uncertain.
<i>Nature of parameters</i>	
Parameters are not repeatable or random.	Parameters are unknown.
They are therefore not random variables, but fixed (unknown) quantities.	They are therefore random variables.
<i>Nature of inference</i>	
Does not (although it appears to) make statements about parameters.	Makes direct probability statements about parameters.
Interpreted in terms of long-run repetition.	Interpreted in terms of evidence from the observed data.
<i>Example</i>	
"We reject this hypothesis at the 5% level of significance."	"The probability that this hypothesis is true is 0.05."
In 5% of samples where the hypothesis is true it will be rejected (but nothing is stated about this sample).	The statement applies on the basis of <i>this</i> sample (as a degree of belief).

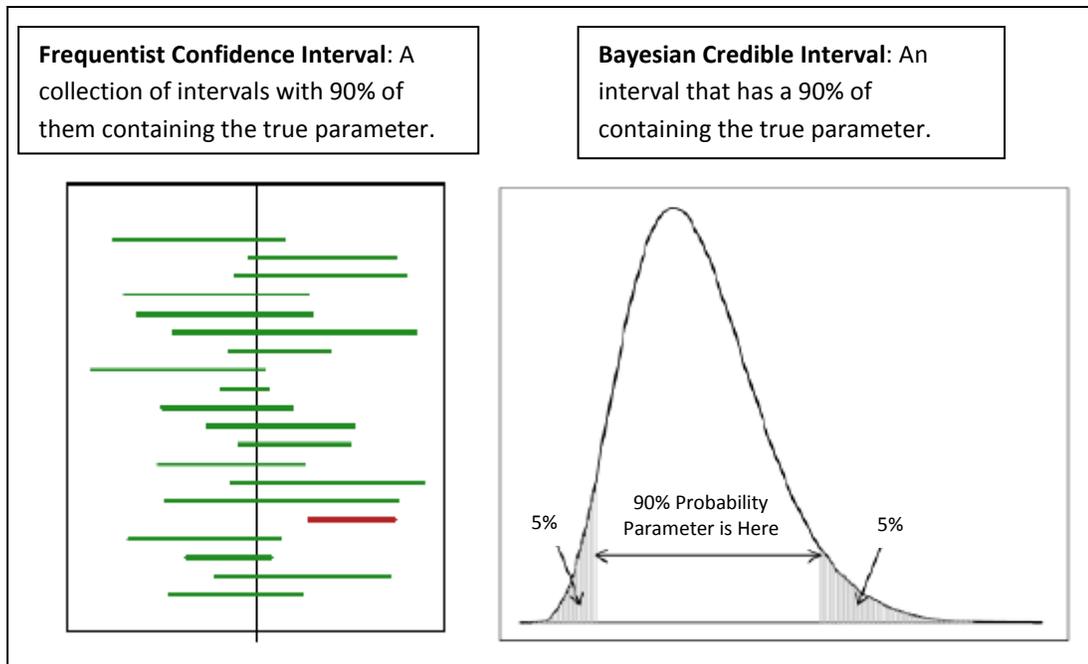


Figure 2-1 Confidence Interval versus Credible Interval (Recchia, 2012)

From this it is clear there are key mathematical differences between the frequentist and Bayesian approach, which has caused severe disagreement. Nevertheless, although there has been much controversy behind the underlying mathematical inferences, the goal of this research is to develop a methodology useful for industrial application. Bayesian statistics has been effective in many practical applications (these will be discussed further below) and can be very useful in demand forecasting for the described supply chain.

The remainder of this section will discuss Bayes' Rule and methods to compute Bayes' Rule. The objective of the Bayesian model is to evaluate posteriors that allow us to calculate an unknown statistic based on a likelihood function and a specified prior distribution. Let's suppose the statistic of interest is demand rate (θ), which is a random and unknown parameter. The prior distribution $p(\theta)$ represents any prior knowledge the modeler knows about demand rate before observing demand. This distribution can be "non-informative" and objective where

the prior does not favor any value of θ , or it can be “informative” and subjective where the prior is formed based on expert opinion. The likelihood function $p(y|\theta)$ is the name given to the model applied to the observed data. The posterior $p(\theta|y)$ is the end goal that allows us to calculate what the modeler believes the true demand rate is based on prior knowledge (prior distribution) and observed data (likelihood function). The prior, likelihood, and posterior are all related in Bayes’ rule:

$$p(\theta|y) = \frac{p(y|\theta)p(\theta)}{p(y)} = \frac{p(y|\theta)p(\theta)}{\int p(y|\theta')p(\theta')d\theta'} \quad (2.1)$$

where the second step transforms $p(y)$ by the law of total probability. Unfortunately, the integral in the denominator is often intractable making Bayesian analysis difficult; however, there are methods to avoid this integral, which will now be discussed.

One approach to avoid the integral in the denominator of Bayes’ rule is to use conjugate priors. The distribution of the likelihood function is usually clear from the data. The key is to assign the right prior for the likelihood function to circumvent the integral when computing the posterior. A conjugate prior is conjugate if the posterior belongs to the same family of distributions as the prior. This makes the computation of the posterior mathematically convenient. The parameters of the prior distribution are called prior hyperparameters, and parameters of the posterior distribution are called posterior hyperparameters. Also, if the likelihood model is in the exponential family, then a conjugate prior exists.

Second, one can use a hierarchical prior. When using a conjugate prior, the prior hyperparameters must be specified and fixed. A hierarchical prior can be conjugate but used to form a probability distribution around the prior hyperparameters. This is called a hyperprior and allows researchers to express uncertainty in the prior assumption. For example, the

modeler believes the prior follows an exponential distribution; however, he is uncertain what the exact parameters should be. This method allows for prior parameters to be considered random (not fixed) and contain uncertainty, albeit at the risk of additional modeling complexity.

Third, Monte Carlo simulation can be used to overcome the intractability of the posterior. The most common Monte Carlo simulations are Metropolis Hastings algorithm and Gibbs Sampler, which both fall under Markov chain Monte Carlo simulations (Carlo, 2004). In these methods one can evaluate what model the likelihood function follows based on observed data. When using this likelihood model in Bayes' Rule, it is likely that the modeler will not be able to obtain the analytical expression of the posterior. However, a computational method, Monte Carlo simulation, for simulating draws from the posterior can be used to estimate the posterior parameters. With a large enough sample of draws, the posterior parameters can be found with a level of arbitrary precision. The level of precision is a function of replications, so it is important to calculate the accuracy of the point estimator in order to reduce estimation error. It is important to note that in classical simulations, the parameters are fixed, so one is not able to incorporate uncertainty like in Bayesian analysis.

In conclusion, the Bayesian paradigm differs from frequentist statistics, which has been the basis of controversy when using Bayes' rule. However, it has been applied successfully in practice, especially for problems with limited data. Bayes' rule involves the relationship between the posterior (the distribution for the unknown parameter), likelihood function (model of observed data), and prior (prior knowledge of the unknown parameter). Computing the posterior can be difficult or impossible. Hence, there are approaches that can avoid the complexity of the denominator in Bayes' Rule such as conjugate priors, hierarchical priors, and Monte Carlo simulation.

2.4.2 Review of Bayesian Literature

The Bayesian approach to forecast demand for spare parts is not new in inventory control. In fact, Bayesian updating has been an active area of research in inventory control literature since the early 1950s. This section will discuss previous literature that relates to demand forecasting using Bayes' Rule. Early literature will be discussed first that focuses on improving inventory models with an assumed Bayesian demand distribution. Then, more recent literature will be discussed that focuses solely on the forecasting component of inventory control and use simulation to compute demand. These methods allow the demand distribution to change based on observed data.

Scarf (1959) was one of the first to propose Bayesian estimation in the context of a periodic review inventory model. He discusses a method to compute optimal inventory levels in the case where the demand distribution contains an unknown parameter and could be described as a Bayesian distribution. He concludes that if the demand distribution of the prior and likelihood belongs to the exponential family, the optimal inventory level is a function of two variables: current stock and past demand. However, functions of two variables are difficult to compute recursively. Scarf (1960) illustrates that if several assumptions are made, it is possible to evaluate optimal inventory levels by recursive computation involving only one variable, known demand distribution. Many who are concerned with deriving the equations of inventory optimality have extended Scarf's work (Hayes, 1969; Iglehart, 1964; Wang and Mersereau, 2013). However, most of this work evaluates inventory levels in an infinite horizon.

In order to optimize inventory in a finite horizon, dynamic programming procedures are used in Zacks (1969), Kaplan (1976), Kaplan (1988), Brown and Rogers (1973). In this work the Bayesian demand models assume Poisson demand, and the conjugate prior for Poisson demand

is a gamma distribution. Of this work in dynamic programming, Brown and Rogers (1973) research is of interest because of their research in Navy aircraft spare parts. They apply their procedure to evaluate optimal inventory requirements on the F-14 program and present some interesting conclusions that can be especially useful for inventory management in the aerospace industry. The most emphasized conclusion is to accept low system reliability early in the life of a program and procure parts as needed (unless vital operational requirements advise otherwise) until sufficient demand information is accumulated. This strategy will significantly minimize cost and waste. Kaplan (1988) evaluated inventory levels for repair parts of new weapons systems and also concluded that buying less until information is accumulated (also known as hedging) is a practical strategy. This emphasizes the need to learn from data before making huge inventory investments.

Comparative studies have been performed that support the use of Bayesian procedures (Azoury and Miller, 1984; Hill, 1999; Soliman et al., 2006). Hill (1999) strives to advocate the use of Bayesian techniques when making decisions with limited data. When there is no information to develop a meaningful prior, a non-informative prior must be used. It can be argued that the Bayesian approach will be least successful when this type of prior is used. Hill (1999) uses a limiting form of the conjugate prior that has an infinite mean, which provides no useful locational information. He compares this with a classical 'point-estimate' approach. His analysis supports the use of Bayesian methods with or without a meaningful prior.

Most of the parametric work up until Hill (1999) assumes Poisson demand and a gamma prior, which is the conjugate prior. One exception is the work of Petrović et al. (1989) where the assumed demand distribution in a certain period follows a binomial (N, p) distribution. N is the known number of opportunities for a single demand to occur in a period with probability p each

time. The conjugate prior for parameter p is a beta distribution. With this assumption, N has to be large for the model to be reasonable. This assumption has proven successful for cases with short request history in work following this research (Dolgui and Pashkevich, 2008; Grange, 1998). Additionally, thus far, dynamic programming has been used to evaluate stock levels in a finite horizon. However, these formulations are typically computationally heavy and the assumptions reduce the model to a single-period model. Further research has been conducted to create simpler models that are able to solve longer planning horizons and non-stationary demand (Kamath and Pakkala, 2002). However, the main focus of this research is to utilize Bayes' rule to forecast demand without changing how stock levels are computed. The remainder of the literature discussed focuses on the demand forecasting component of inventory control.

Aronis et al. (2004) utilizes the Bayesian approach to forecast spare parts demands for electronic equipment. The goal is to propose an improved yet practical method without altering the inventory control policy. They perform two case studies. First, they determine the best method to select prior parameters. They assume demand follows a Poisson model, and the conjugate prior for this model is a gamma (α, β) distribution. In order to find the two prior parameters (α, β), a system of two equations is needed. In practice it is common to use the mean or mode of the demand rate (equation 1) with the percentile of the distribution (equation 2). To determine the percentile of the distribution experts specify X of the percentile. In other words, in 95% of the cases, the failure rate does not exceed X times the mean or mode. Four alternative methods using the mean or mode with varying X values are studied. All approaches return stock levels plus or minus one of each other. This analysis proves that the end result is not sensitive to the prior because the prior importance diminishes quickly. Second, based on this conclusion they propose a new method that includes the option to weight the mean

demand rate in order for the prior to hold importance in the equation. To compute prior parameters they use the mean and percentile method. A case study was performed to determine the optimal weight and X values for this approach. Based on these optimal values, the case study compared the inventory stock levels of the proposed Bayesian method and the current method. The proposed approach resulted in lower stock levels at a 95% service level. However, the authors note that the case study evaluates the predicted service levels and not the true service levels.

More recent work uses simulation to solve Bayes' rule in order to allow the demand distribution to vary depending on observed data. Gelfand and Smith (1990) explore three Monte Carlo sampling approaches (Stochastic substitution, the Gibbs sample, and the sampling-important-resampling algorithm) to calculate the Bayesian posterior density. They conclude substitution or Gibbs methods consistently provide better performance in terms of efficiency. These methods have been used in latter work to compute the Bayesian posterior, and due to increased advances in computing power, recent literature uses these Monte Carlo approaches to forecast demand. However, literature on utilizing simulation to forecast demand is sparse because it is mostly used to compare forecasting methods. It is also worth mentioning that the bootstrap model of Willemain et al. (2004) uses simulation to forecast demand. However, the parameters are fixed, so it is unable to capture parameter uncertainty like in Bayesian statistics.

Yelland (2010) uses a Bayesian Gibbs Sampler approach to estimate part demand for a major vendor of network computer products. Hierarchical priors are used to pool demand patterns for parts with historical records (or established parts) to produce initial parameter estimates for parts with little to no demand. This procedure can also be found in Duncan et al. (2001). They perform a case study to test forecast performance between judgmental forecasts,

exponential smoothing, and the Bayesian model. The Bayesian model performed best in all forecast accuracy tests. However, the authors are still wary of recommending the proposed method in application because it is complex and the MCMC algorithm is computationally expensive to run.

Muñoz and Muñoz (2011) propose two different Bayesian simulation methods to forecast the demand of spare parts. The first method uses a posterior sampling approach to estimate the mean. However, this method requires a complex algorithm to generate samples from the posterior, which can be difficult to analytically express. The second method uses the Independence Sampler Markov Chain Monte Carlo (MCMC) algorithm. The independence sampler is a special case of the Metropolis-Hastings algorithm and avoids the complex algorithm in the posterior sampling method. After developing the methods, the paper applies these models to a car dealership's dataset in order to illustrate how to utilize these methods when forecasting the demand for spare parts. They conclude three things. First, the accuracy of the point is dependent on number of simulation replications, so it is important to determine number of replications based on how accurate one wants the point estimator to be. Second, the results show that the prior is dominated by the data extremely quick. They then simulated a sample of data with a higher failure rate, and the prior was more influential. Third, they emphasize the importance of the models ability to represent a real system (such as parameters and distributions).

Rahman and Sarker (2012) explore the Bayesian approach to forecast intermittent demand for seasonal products. In the Bayesian model, the Gibbs sampling algorithm is used to compute the parameters and forecast intervals. The demand structure is found by the SARIMA model, and the prior is a non-informative prior similar to work in Gelman et al. (2013) and

(Congdon, 2003). Rahman and Sarker (2012) compare the Bayesian method to SARIMA and multiplicative exponential smoothing, which are both prominent methods to forecast seasonal products. Several forecast error metrics and a cost factor method computed by dynamic programming are used to compare three forecasting methods. The results illustrate that the Bayesian method performs superior to the other models and is effective at forecasting seasonal and intermittent demand. De Alba and Mendoza (2007) also use a Bayesian approach to forecast seasonal demands. They conclude that this method is the most effective for short time series, which they define as less than 2-3 years. However, they do not look into intermittent demand.

Overall, the Bayesian framework has been applied extensively in the field of forecasting spare part's demand. Most of the early work assumes a Poisson demand distribution with a gamma prior (a conjugate prior). With this assumption the work revolved around characteristics of optimal solutions in the periodic review case, or it explored dynamic forecasting methods to optimize inventory in a finite time horizon. Aronis et al. (2004) was unique in the sense that it used this same demand assumption; however, it applied the demand forecasts to the company's current inventory policy. More recent literature uses Monte Carlo simulation to forecast demand in order to solve complex problems. It has proven effective; however, running these algorithms can be expensive, and research in this area is small. Further empirical results are needed.

2.5 Critique of Literature

The intent of this section is to critique the reviewed literature in order to find gaps in the field of demand forecasting for spare parts. This discussion will lead into how this research intends to bridge these gaps and ultimately, add to existing literature. A critique on demand

forecasting literature as a whole is discussed first. Then, the three forecasting techniques reviewed are critiqued: judgmental, statistical, and Bayesian.

First, demand forecasting literature as a whole is critiqued. The literature on demand forecasting for spare parts has recently gained renewed interest due to concerns in maintenance modeling and intermittent demand. This has advanced the theoretical knowledge in spare parts forecasting; however, current literature has expressed a gap between research and practice in this field. Most of the methods developed for spare parts forecasting are commonly neglected in practice due to data unavailability, model complexity, and forecasting support systems. This illustrates motivation to develop a practical demand forecasting recipe that can be used in industry. Hence, this research will propose a practical method that can be easily integrated in practice. Additionally, there are no conclusive results on which forecasting method is best. The majority of comparative studies use forecast accuracy metrics that are not suitable for data with many zeros, and this can result in biased conclusions. Also, many of these studies compute inventory with a specified fill rate goal and conclude that a method is best because the inventory investment is reduced at the same specified fill rate goal. However, in reality, the “true” fill rate performance will vary from the specified fill rate goal used in optimization. Thus, this research will evaluate the impact of the different methods on “true” fill rate performance using simulation and determine if these results agree with the forecast accuracy and inventory optimization (evaluated inventory costs with specified fill rate) analysis.

Next, the three forecasting methods reviewed in this research are critiqued. Judgmental forecasting is reviewed first. These forecasts are formed by expert opinion and are commonly used in practice as an adjustment factor for statistical methods. Benefits of judgmental adjustments for slow-moving and intermittent items are shown in Syntetos, Nikolopoulos, et al.

(2009). Thus, judgmental input along with statistical methods can be extremely effective. However, manual adjustments are extremely difficult when managing a large number of parts. Hence, when a large number of items exist, a systematic approach to include expert opinion is almost necessary. This research proposes a methodology that includes judgmental input in the estimates of demand. Additionally, there is limited literature that expresses the impact of judgmental forecasting in the context of demand forecasting for SKUs. Further research is needed on implications of forecasting adjustments on prediction and inventory (Boylan and Syntetos, 2009). This research will address this by illustrating the merits and consequences of judgmental forecasting on forecast accuracy, inventory costs, and service levels.

Four statistical forecasting techniques are reviewed: time series, causal factors, bootstrap, and neural networks. Traditional time series methods such as moving average and simple exponential smoothing are commonly used in industry due to practicality. However, if a trend exists, these methods lag in forecast. This is especially concerning when new aircraft programs steadily increase in fleet size (thus, total operating hours). Hence, these traditional time series approaches are not recommended for new aircraft programs, and forecasting for new programs is in the scope of this research. In the case of initial life cycle forecasting, literature recommends causal models (Boylan and Syntetos, 2008).

Statistical forecasting methods have also been developed to address intermittent demand, which is a characteristic of most spare parts. Croston's method, a time series approach, is known for its' ability to forecast intermittent demand and is integrated in some of the leading software packages. However, this method assumes the functional form of demand. Therefore, bootstrapping and neural network approaches can be used to estimate functional relationships of intermittent demand. Bootstrapping is not appropriate for this thesis because it gives

demands in a lead-time. This type of information is not easy to plug into inventory optimization software that optimizes inventory based on fill rates, and this research has not found a bootstrap method that is capable of computing demand based on fill rates. Also, neural network approaches can be are complex and very specific to the given problem making these unattractive for this thesis.

After critiquing the statistical forecasting literature, causal models seem the most appropriate for new aircraft programs, and Croston's method or its' variants seem the most appropriate for intermittent demand. However, these methods along with the other statistical methods assume past data represents future behavior. However, initial demand can vary from long term experience; therefore, forecasts based on initial demand can lead to extremely poor performance. In order to cope with this, many use judgmental forecasts until a "significant" amount of demand is accumulated. However, many struggle with the question of when there is enough data to make this transition from judgmental forecasts to statistical methods, which can extremely impact service levels. This thesis illustrates how this transition can negatively impact service levels, and the proposed method eliminates the need to make this transition.

Bayesian forecasting provides a systemic way to update prior information (such as expert opinion) as observed demand occurs. Bayes' rule can be extremely difficult to solve, so three methods were discussed to avoid intractability: conjugate priors, hierarchical priors that are conjugate, and Monte Carlo simulation. Most of the early work in Bayesian literature assumes Poisson demand (likelihood function) with a gamma conjugate prior. With this assumption, the research focused on inventory optimization characteristics. Aronis et al. (2004) was unique in the sense that it used this same demand assumption; however, it applied the demand forecasts to the company's current inventory policy. More recent literature uses Monte Carlo simulation

to solve more complex problems that make Bayes' rule intractable. It has proven effective; however, running these algorithms can be expensive and research in this area is small.

The model proposed in this thesis uses Bayes' rule to forecast demand. It is important to note that this thesis will propose a forecasting method without changing the inventory policy, similar to the approach of Aronis et al. (2004). Furthermore, this research is focused on practicality constraints, and it is common in practice for companies to aggregate data in a way that does not support the calculation of a demand distribution. Hence, a demand assumption must be made. Most of the parametric work assumes the likelihood function (observed demands per time) follows a Poisson distribution. However, for this research, the statistic of interest is time between demands. If demands per time are Poisson, then the time between demands is exponentially distributed. Therefore, this research will assume that the likelihood function (observed time between demands) is exponentially distributed. Further, an informative conjugate prior is used to circumvent the denominator in Bayes' rule. This approach is extensively used in literature and is practical for application. The conjugate prior for an exponential likelihood function is a gamma(α , β) distribution. The methods in the literature that compute prior parameters are not supported by the case study data, so this research will assume the prior distribution is also exponentially distributed (gamma distribution with $\alpha=1$ and $\beta=\text{mean}$). Also, assuming the prior time between failures is exponentially distributed seems more appropriate than assuming the prior demands per time follow a gamma distribution, which is the prior used in most literature. However, when making this exponential prior assumption, if no demand has occurred, the computed mean time between failures is extremely high. Thus, this research proposes two business rules to more accurately represent demand predictions in this situation. These rules have not been presented in the previous Bayesian literature.

In conclusion, this section examines prior research to highlight benefits and shortcomings of existing approaches. Aside from the methods, a gap between research and practice has also been pointed out in the literature, motivating the need for a practical solution that can be easily integrated into industry. Also, most of the comparative studies use forecast accuracy or inventory optimization (evaluated inventory costs with specified fill rate) results to compare methods (Bacchetti and Saccani, 2012). However, these results can lead to biased conclusions. Thus, this thesis will evaluate the impact of the forecasting methods on “true” fill rate performance using simulation and determine whether these results agree with the forecast accuracy and inventory optimization analysis. Second, judgmental forecasting is critiqued. There is limited work in judgmental forecasting for SKUs. Thus, this thesis will illustrate the merits and consequences of this approach on forecast accuracy, inventory costs, and service levels. Third, statistical forecasting is critiqued. Statistical methods assume past data represents future behavior. However, when parts have limited data due its’ new or slow-moving nature, past data can lead to poor performance. This thesis incorporates prior information to address this initial variation. In practice, companies commonly try to overcome this by using judgmental data until a “significant” amount of demand is accumulated, and this research will show how the transition from judgmental to statistical forecasting can negatively impact service levels. Fourth, Bayesian forecasting is reviewed. The proposed model uses Bayes’ rule to forecast demand in a manner similar to Aronis et al. (2004) because it does not change the inventory optimization model, and weights are applied to prior parameters. The weights give this research the ability to incorporate one’s confidence in engineering estimates, and this application of using prior weights differs from previous literature. Furthermore, this research is unique to Bayesian research in the following ways. The proposed model assumes the likelihood function is exponentially distributed where most work assumes the likelihood model is Poisson.

Additionally, this research develops a method to more accurately depict demand when zero demands have occurred. This method has not been found in preceding Bayesian forecasting literature. In conclusion, these are all reasons why this research adds to existing research.

Chapter 3 – Problem Description and Model Formulation

3.1 Problem Overview

Demand estimation is an essential component in the analysis of inventory systems. In practice, demand distribution parameters are commonly fixed subjectively using expert opinion or statistically estimated using historical demand. However, it is almost impossible to exactly evaluate the true values of these parameters especially for new aircraft programs with an absence of demand data. Additionally, when a new program is employed, demand may experience variation or exhibit a trend, and inventory optimization can be very sensitive to changes in demand rate. To overcome the difficulty of limited data the Bayesian framework can be used, which is the proposed method in this research. The Bayesian methodology used updates initial estimates of demand formed by experts as observed demand occurs. The overall goal of this research is to utilize the Bayesian framework to forecast demand without changing how inventory is optimized. This chapter will show how demand is currently estimated and formulate the proposed Bayesian forecasting model.

3.2 Current Demand Forecasting Practices

This section discusses how the aircraft manufacturer (which the case study is based upon) currently forecasts demand for spare parts. This method will be referred as the Current Method throughout this thesis. The initial estimates used when no demand information is available are discussed first. Then, the rules used to transition from these initial estimates of demand to observed demand are described. Last, the statistical forecasting method applied to actual demand is explained.

Engineering estimates are initially used to forecast demand. These estimates approximate the mean time between failures and can come from a variety of sources. A

program manager’s confidence varies based on the source of the engineering estimate, and there is currently no systematic approach to incorporate their confidence in predicted demand.

Table 3.1 groups parts based on how confident managers are in the engineering estimate relative to the specific program used in the case study.

Table 3.1 Level of Confidence in Engineering Estimate for Different Parts Groups

Part Group	Description of Part Group	Confidence in Engineering Estimate (0-100% where a higher value means more confident)
Standard Parts	Parts that are standard to all aircraft (i.e. nuts, bolts, etc.)	100%
Commercial Common Part – Part <u>is</u> consistent with commercial utilization	Part is common with commercial fleet and utilization is consistent with commercial fleet	95%
Commercial Common Part – Part <u>is not</u> consistent with commercial utilization	Part is common with commercial fleet but utilization is not consistent with commercial fleet	75%
Off the Shelf Part	Part is common to platform and utilization	95%
Unique Part	Design is unique to this program	25%
Unknown Part	Insufficient information available to categorize part	25%

After establishing the engineering estimates, managers must decide when the estimates of demand should transition from engineering estimates to observed demand. Transitioning to actual demand can be triggered by:

1. Accumulation of a certain amount of operating hours
2. Accumulation of a certain amount of demand

The approach used for the aircraft of interest utilizes the second method. The transitioning rule is:

The engineering estimate will be used for 2 years. Then, after 3 demands in a rolling 12 month period has occurred, estimates of demand will transition from engineering estimates to actual demand. After this transition engineering estimates will no longer be used.

This transitioning point could result in increased cost and waste. Suppose the engineering estimate predicts a part is a high demand part; however, in reality, it is a low demand part. The part might not accumulate enough demand to switch over to observed demand. Thus, a significant amount of unused capital in inventory will result. Also, the engineering estimate can be a poor estimate of a part's true behavior, and two years is a long time to optimize inventory based on poor estimates of demand. Hence, an approach that includes the true behavior of a part early in the life of program is needed.

Once parts transition to observed demand, a blend of simple exponential smoothing and causal factors are utilized to calculate the number of demands in the next period. The forecaster specifies how many periods to forecast. The equations used to compute the blended forecast are:

Blended Forecast:

$$d_{t+1} = \gamma(SES_{t+1}) + (1 - \gamma)(CF_{t+1}) \quad (3-1)$$

Simple Exponential Smoothing:

$$SES_{t+1} = \alpha d_t + (1 - \alpha)d_{t-1} \quad (3-2)$$

Causal Factors:

$$CF_{t+1} = \frac{OH_{t+1}}{\sum_0^t OH / \sum_0^t d} \quad (3-3)$$

Table 3.2 defines the notation in equations 3-1, 3-2, and 3-3 and shows the values of the parameters if they are constant.

Table 3.2 Blended Forecast Equation Notation

Parameter	Equation	Definition	Value if constant
t	3-1,3-2,3-3	Current time period	
d	3-1,3-2,3-3	Number of demands	
γ	3-1	Blending factor	t ≤ 6 then $\gamma=.25$ t > 6 then $\gamma=.5$
SES	3-1, 3-2	Simple exponential smoothing forecast for number of demand	
CF	3-1, 3-3	Causal factors forecast for number of demand	
α	3-2	Simple exponential smoothing constant	.10
OH	3-3	Operating hours	

As discussed in Chapter 2, the blended forecast has some concerns when forecasting low and intermittent demand. The method assumes past data represents future demand. However, early behavior of a part can misjudge future demand. Also, SES is known to lag when there is a trend. Thus, SES might perform inferior to other methods capable of handling trends.

In conclusion, this section summarizes the current practices used to forecast demand. When demand is not available, engineering estimates that approximate the mean time between failures are used. These estimates come from a variety of sources; however, there is currently no systematic method to account for how reliable these estimates are. Next, after 3 demands in a rolling 12 month period have occurred, estimates of demand will transition from engineering estimates to actual demand. This method could result in exaggerated waste when the engineering estimate predicts the part to be a high demand part, and in reality it is a low demand part. Also, if the engineering estimate is a poor estimate of demand, two years is a long period to optimize stock levels based on poor estimates. Furthermore, when actual demand is used, a blending of simple exponential smoothing and causal factors are applied to the data to approximate demand rate. However, this approach assumes past behavior represents future

demand, and this can assumption can extremely effect estimates of demand early in the life of a program.

3.3 Formulating Demand Forecasting Using Bayes' Rule Equation

Bayesian forecasting responds to the critiques of the Current Method (i.e. ability to handle increasing operating hours, learn from observed demand immediately, and include one's confidence in the engineering estimates). This section will apply Bayes' rule to demand forecasting. Based on the aggregation of case study data, the model assumptions and methods used to avoid intractability of Bayes' rule are discussed first. Then, the formulation of Bayes' rule to forecast demand is presented. This will include how to calculate the credible interval and explain the developed business rules when zero demand exists.

3.3.1 Model Assumptions

The overall goal of this sub-section is to evaluate the assumptions needed to formulate Bayes' rule in this research. The aggregation of data used in the case study is illustrated first. This will lead into discussion on the assumptions used to formulate the Bayesian forecasting model.

First, it is important to illustrate how the demand data used in the case study is collected. The aircraft manufacturer collects part demand dates and fleet operating hours per month. An example of this information is shown in Table 3.3 and Table 3.4, respectively.

Table 3.3 Part Demand Dates

Part Number	Order Quantity	Order Date
Part 1	1	7/11/2003
Part 1	2	10/14/2006
Part 2	3	7/14/2003
Part 3	1	1/2/2002

Table 3.4 Aggregated Flight Hours per Month

Part Number	5/1/2001	6/1/2001	7/1/2001
Part 1	50	100	250
Part 2	100	200	500
Part 3	25	50	125

This aggregation does not support the calculation of a demand distribution. It only allows for the calculation of a single parameter of an assumed demand model (i.e. average operating hours per failure). In order to calculate the demand distribution, the data collected must show the total operating hours since the last demand for each individual asset on all aircrafts. This will be further discussed in Chapter 6. However, the data is not collected in this manner, so a demand assumption must be made.

Most of the parametric work assumes the likelihood function (observed demands per operating hour) follows a Poisson distribution. However, for this research, the statistic of interest is operating hours per demand because it is more intuitive for low demand parts than demands per operating hour. If demands per operating hour follow a Poisson distribution, then the operating hours per demand is exponentially distributed. Therefore, this research will assume that the likelihood function is exponentially distributed.

Conjugate priors will be used to avoid the integral in Bayes' rule. These types of priors are algebraically convenient and utilized in most Bayesian parametric work. Hierarchical priors add additional complexity, and the company would have to purchase software designed for Bayesian analysis to compute hierarchical priors. Also, the data collected for the case study does not support the calculation of a demand distribution, so there is no need to utilize Monte Carlo simulation. Furthermore, the conjugate prior for an exponentially distributed likelihood

function is a gamma distribution, and this prior will be considered informative. Experts establish engineering estimates (mean time between failures) before aircrafts are deployed, and this information is easily accessible.

After establishing the prior assumption, the prior parameters (α , β) must be evaluated. Aronis et al. (2004) use a system of equations with mean or mode and percentile of the distribution. Recchia (2012) recommends a system of equations with mean and variance. However, the percentile of the distribution and variance are not available. Hence, this research will assume the prior distribution is also exponentially distributed. An exponential distribution is a special case of the gamma distribution ($\alpha=1$, $\beta=\text{mean}$), so this assumption is valid. Also, although Poisson demands per operating hour (likelihood assumption in most research) equates to an exponentially distributed operating hours per demand, assuming the prior operating hours per demand is exponentially distributed (as assumed in this research) seems more logical than assuming the prior demands per operating hour is gamma (as assumed in most parametric Bayesian research).

3.3.2 Formulating Bayesian Model

Given the demand data is aggregated and does not support the evaluation of a demand distribution, the formulation of the Bayesian model requires the following assumptions:

1. **Likelihood Function:** Observed mean time between demands is exponentially distributed.
2. **Prior:** Engineering estimates (mean time between failures) are exponentially distributed. However, the prior is illustrated as a gamma function. This is appropriate because the exponential distribution is a special case of the gamma ($\alpha=1$, $\beta=\text{mean}$) distribution.

Based on these assumptions, the posterior will be formulated using Bayes' rule. The posterior will be used to evaluate operating hours per demand (or mean time between failures). The unknown parameter of interest is λ , which is defined as operating hours per demand.

Likelihood function (exponential):

$$L(n|\lambda) = \lambda^n e^{-\lambda \sum_{i=1}^n x_i} \quad (3-4)$$

Prior (gamma):

$$g(\lambda; r, v) = \frac{v^r \lambda^{r-1} e^{-v\lambda}}{\Gamma(r)} \propto \lambda^{r-1} e^{-v\lambda} \quad (3-5)$$

v^r and $\Gamma(r)$ remain constant in respect to λ , so these parameters can be ignored when computing λ .

Posterior (gamma):

$$p(\lambda|n) = \frac{p(n|\lambda)p(\lambda)}{p(n)} = \frac{p(n|\lambda)p(\lambda)}{\int p(n|\lambda')p(\lambda')d\lambda'} \propto p(n|\lambda)p(\lambda) \quad (3-6)$$

The integral (or partition function) in the denominator stays constant with respect to λ , so it can be ignored when computing λ . The posterior is computed below:

$$\begin{aligned} p(\lambda|x, r, v) &\propto p(n|\lambda)p(\lambda) & (3-7) \\ &\propto (\lambda^n e^{-\lambda \sum_{i=1}^n x_i})(\lambda^{r-1} e^{-v\lambda}) \\ &\propto \lambda^{n+r-1} e^{-(v+\sum_{i=1}^n x_i)\lambda} \end{aligned}$$

Equation 3-7 is in the form of a gamma distribution and is equivalent to $\text{Gamma}(r + n, v + \sum_{i=1}^n x_i)$. It will also be expressed as $\text{Gamma}(r', v')$. The full posterior equation follows:

$$p(\lambda|x, r, v) = \frac{v + \sum_{i=1}^n x_i}{\Gamma(r)} \lambda^{r+v-1} e^{-(v+\sum_{i=1}^n x_i)\lambda} \quad (3-8)$$

After formulating the posterior, the mean time between demands can be computed. This is shown in equation 3-9.

Mean Time between Demands:

$$\frac{r'}{v'} = \frac{r+n}{v+\sum_{i=1}^n x_i} \quad (3-9)$$

Notation for equation 3-9:

r: engineering estimate (mean time between failure)

n: the number of operating hours in observed data

v: 1 (the exponential distribution is a special case of gamma when r=1 and v=mean)

x: the number of demands in observed data.

Additionally, previous Bayesian research expresses that the prior importance in the posterior diminishes quickly by actual demand data. Due to this, equation 3-10 shows the updated demand calculation where a weight (w) is applied to prior parameters so that the prior can maintain importance. This weight (w) parameter is also used to allow managers to incorporate their initial confidence in the engineering estimates. The new mean time between failures can be calculated in equations 3-10 and 3-11.

Posterior: $\text{Gamma}(r',v') = \text{Gamma}(w * r + n, w * v + \sum_{i=1}^n x_i)$ (3-10)

Mean Time between Failures: $\frac{r'}{v'} = \frac{w*r+n}{w*v+\sum_{i=1}^n x_i}$ (3-11)

These weights (w) are computed based on the manager's confidence levels in the engineering estimates shown in Table 3.1. The weight increases prior parameters by the confidence level given in this table. For example, if the manager is 95 percent confident in the engineering

estimate, the prior parameters for that part is increased by 95 percent, which is equivalent to a weight (w) value of 1.95.

After the posterior is computed, a credible interval can be evaluated. This interval shows the probable range of values that the “true” mean time between failure lies. This can be useful when evaluating uncertainty in the posterior mean or determining whether the engineering estimate is a “good” estimate of the mean. The credible interval can easily be computed in Excel using the Gamma Inverse function. For example, if the specified confidence level is 95%, the excel function for the lower quantile is Gamma.Inverse(.025, r', v'), and the excel function for the upper quantile is Gamma.Inverse(.975, r', v').

The Bayesian model to predict demand (posterior equation) and demand uncertainty (credible interval) has been presented. However, when no observed demands have occurred, the Bayesian posterior in this method is extremely high. This is best be shown by example. Suppose the prior engineering estimate is 500 operating hours per failure with a prior weight of 1.5, and the engineers have observed 0 demands in 700 operating hours. The calculation for Bayes’ estimate of mean time between failures (Bayes’ MTBF) would be:

$$\text{Bayes' MTBF: } \frac{r'}{v'} = \frac{w*r+n}{w*v+\sum_{i=1}^n x_i} = \frac{1.5*500+700}{1.5*1+0} = \mathbf{966.67 \text{ operating hours / failure}}$$

967 operating hours per failure is larger than the engineering estimate (500 operating hours per failure) and observed demand (700 operating hours with no failures). This does not accurately depict how parts are behaving, so two business rules are created below.

First, when a part has experienced zero failures and the observed operating hours are less than or equal to the engineering estimate, the observed data does not oppose the engineering estimate. Hence, the engineering estimate can be used as the estimate of demand.

For example, suppose the engineering estimate is 500 operating hours per failure, and the engineers have observed zero demands in 300 operating hours. The observation of zero demands in 300 operating hours has not reached the engineering estimate of 1 failure in 500 operating hours. Thus, this information does not oppose the engineering estimate, and the engineering estimate can be used to predict demand.

Second, when a part has experienced zero failures and the observed operating hours are greater than the engineering estimate, the observed data opposes the engineering estimate. Hence, this research includes a second business rule that allows the estimate of demand to incorporate this knowledge. When this scenario occurs, one can input 1 failure per the observed amount of operating hours in the likelihood function. For example, suppose the engineering estimate is 500 operating hours per failure with a prior weight of 1.5, and the observed data is zero failures over 700 operating hours. If 1 demand per 700 operating hours is used in the likelihood function, the posterior mean equals 580 operating hours per demand. This allows the estimate of demand to move away from the engineering estimate and towards the observed data even when no demands have occurred.

However, the question is when the model should incorporate these two rules. The prior is exponentially distributed with a mean equal to the engineering estimate, and one can use this knowledge to compute a transition point. The user must first specify a probability, and this study will use 75% as the specified probability. Then, the transition can occur at the computed operating hours per failure where there is a 75% likelihood that the true mean is less than or equal to the computed operating hours per failure (or transition point). To illustrate this, the example above will be used (engineering estimate = 500 operating hours / failures and observed data = 700 operating hours / failure) with a specified probability of 75%. There is a 75%

probability that the true mean is less than or equal to approximately 515 operating hours per failure. Thus, 515 operating hours per failure will be the transitioning point. The observed operating hours is greater than the engineering estimate. Thus, 1 failure (n) and 700 operating hours (x) are inputs in the likelihood function. The two business rules are summarized as follows:

- 1) When actual demand is zero and observed operating hours is **less than** or equal to transition point, then the engineering estimate is used ($n=0$, $x=0$).
- 2) When actual demand is zero and observed operating hours is **greater than** the transition point, then likelihood function will incorporate 1 failure in the amount of observed operating hours ($n=1$, $x=\text{total observed operating hours}$).

In conclusion, this section applies Bayes' rule to develop the proposed forecasting method. Observed demand follows an exponential distribution and the engineering estimate follows a gamma distribution. Based on these assumptions, the posterior is computed resulting in a gamma distribution. Next, a weight parameter is added to the model because the prior distribution can easily become overwhelmed by data. This weight parameter provides the ability to include one's confidence in prior information. Further, after formulating the posterior, the steps to compute the credible interval are discussed, and last, two business rules were created to accurately reflect observed demand when no demand has occurred.

3.4 Problem Description and Model Formulation Summary

This section first explains how the aircraft manufacturer currently predicts demand. This forecasting approach will be referred to as the Current Method in this thesis. The Current Method uses engineering estimates until two years of demand has accumulated. Engineering estimates come from a variety of sources, and there is currently no systematic way to include

one's confidence in these estimates. Also, if the engineering estimates are poor estimates of demand, then two years is a long time to optimize stock levels based upon these estimates. Furthermore, after this two year period and a part has accumulated three or more demands in a rolling twelve month period (the transition point from engineering estimates to observed demand), a blend of simple exponential smoothing and causal factors is used on observed demand. This statistical method assumes that past behavior represents future demand. However, when parts have limited data due its' new or slow-moving nature, past data can lead to poor performance. Further, the transition point of three demands can be especially problematic when the engineering estimate predicts a part to be high demand, but in reality the part is low demand. The engineering estimate will not be revised and unused capital in inventory will result. Additionally, when no demands have occurred and the observed mean time between failures exceeds the engineering estimate, the Current Method will not illustrate this because of the implemented transition point to observed demands.

Next, the formulation of the proposed Bayesian forecasting model is presented. The proposed forecasting method applies Bayes' rule to respond to the concerns of the current forecasting practices. The formulated method incorporates prior information to address initial variation in observed behavior, which eliminates the assumption that only past behavior represents future demand. Engineering estimates are updated as observed demand occurs, and thus, the transition point from engineering estimates to observed demand is removed. Weights are applied to prior parameters so that managers can incorporate their confidence in the engineering estimate, and two business rules are presented when no demand has occurred. This will more accurately reflect the observed mean time between failures in this scenario. Last, it is important to note that the formulated model is a practical solution that addresses common data constraints and can be easily implemented in industry.

Chapter 4 – Case Study Methodology

4.1 Overview of Case Study Methodology

The overall objective of the case study is to validate the proposed Bayesian model for estimating the demand for spare parts. This case study is based upon information provided by a leading aircraft manufacturer company on an aircraft program with 3 ¼ years of data. This program is extremely young relative to its' service life. However, the goal of this thesis is to provide a forecasting method that is especially targeted to support new programs. The program reaches steady state by the second year of operation which implies that all aircraft ordered will be delivered by the second year, and the total operating hours each period remain steady. The forecasting methods compared in the case study are Engineering Estimates, the Current Method, and the Bayesian Method. These methods will be compared in three sections: the forecasting method's impact on (1) forecast accuracy, (2) inventory, and (3) performance. Additionally, parts are grouped in three networks: Line Replaceable Unit (LRU), Consumable, and Maintenance Significant Consumable (MSC). The forecast accuracy analysis investigates parts in all networks while the inventory and performance analysis only consider the LRU network. First, these networks will be explained. Then, the three sections of the case study will be discussed: impact to forecast accuracy, impact to inventory, and impact to performance.

4.2 Overview of Networks

The inventory optimization model (discussed later in this chapter) contains three grouping of parts known internally as networks. Each network is optimized separately to achieve an 80% target fill rate goal. The three networks are:

1. Line Replaceable Units (LRU) Network – Parts that are replaced by repair.

2. Consumables Network – Parts that are not repairable. The old parts are condemned, discarded, and replaced with a new part.
3. Maintenance Significant Consumables (MSC) Network – Consumable parts that are critical to operation of an aircraft.

The majority of parts lie within the LRU network, and currently, most of the invested inventory is in LRU parts. Figure 4-1 illustrates the percentage of parts in each network, and Figure 4-2 illustrates the current percentage of inventory investment in each network. Thus, the simulation model used to evaluate fill rate was configured to only consider parts in the LRU network.

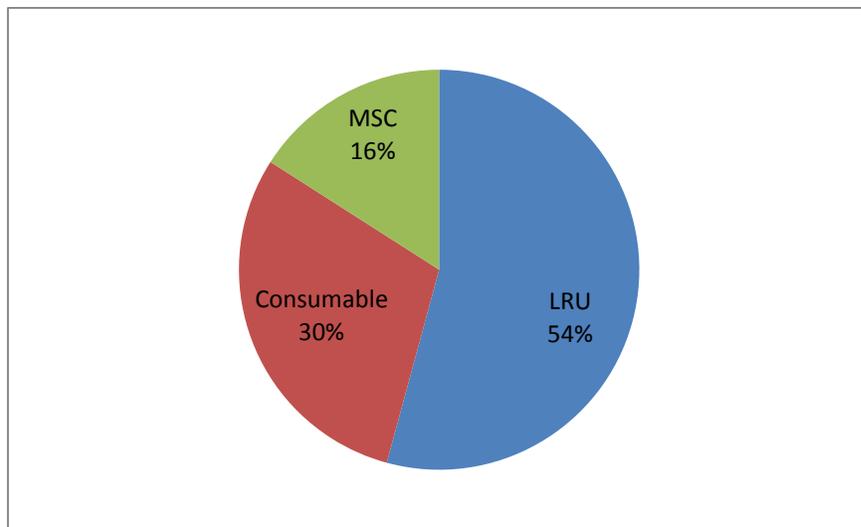


Figure 4-1 Part Breakdown by Network

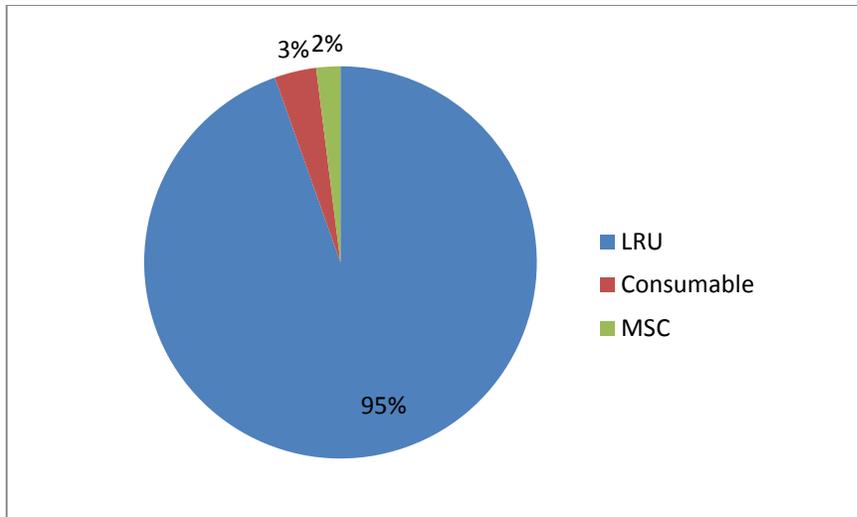


Figure 4-2 Current Inventory Investment by Network

4.3 Impact on Forecast Accuracy

The first portion of the case study determines the value of using Bayes' rule on demand forecast accuracy. The three methods compared are Engineering Estimates, Current Method, and Bayesian Method. Also, this section looks at parts in all networks. Three forecast accuracy metrics are used, and these metrics are explained first. After, the forecast accuracy methodology is discussed. This will show how the forecast accuracy portion of the case study will be executed.

4.3.1 Forecast Accuracy Metrics

The forecasting methods will be compared to observed demand by three forecast accuracy metrics: mean absolute deviation (MAD), mean absolute percentage error (MAPE), and the ratio of MAD to mean (MAD/Mean). The majority of parts are low demand parts, so the time unit used within these metrics is annually.

MAD is the primary metric used for evaluating forecast accuracy. It measures the average deviation of forecasts from actuals and is very easy to compute. MAD is calculated in

equation (4-1) where A_n is the actual demand value, F_n is the forecasted demand value, and N is the number of parts.

$$MAD = \frac{1}{N} \sum_{n=1}^N |A_n - F_n| \quad (4-1)$$

When comparing a method across different time series, MAD can be deceptive. It is a scaled metric and should not be used to compare a method between different series. Thus, MAPE is used to overcome this. It is a scale-free metric and measures the mean absolute error between the forecasts and actuals. For example, if the MAPE is 40%, this tells the forecaster that on average the forecasts overestimate or underestimate the actual demand by 40%. Equation (4-2) is the calculation for MAPE.

$$MAPE = \frac{100}{N} \sum_{n=1}^N |A_n - F_n| \quad (4-2)$$

MAPE is commonly used in industry; however, it can also be a poor estimator of accuracy for certain scenarios. First, MAPE has problems when demand is fluctuating greatly (Kolassa and Schütz, 2007). Second, an underestimated forecast error has an upper bound of 100%; however, an overestimated forecast error has no upper bound. This creates a bias (Armstrong and Collopy, 1992). Last, MAPE cannot be used when there are zero values in actual demand, which is the case for intermittent and low demand parts. In order to cope with these issues, the ratio of MAD to mean (or MAD/Mean) can be used. This metric is comparable across series and can handle intermittent and fluctuating demand series. MAD/Mean is calculated in equation (4-3).

$$\frac{MAD}{Mean} = \frac{\frac{1}{N} \sum_{n=1}^N |A_n - F_n|}{\frac{1}{N} \sum_{n=1}^N A_n} \quad (4-3)$$

MAD, MAPE, and MAD/Mean will all be utilized in the forecast accuracy portion of the case study. MAD and MAPE are commonly used in industry, and MAD/Mean can handle intermittent demand items. Also, it is important to note that this analysis does not compare time series. It compares forecasting methods in each annual series, so MAD is appropriate for this study.

4.3.2 Impact on Forecast Accuracy Methodology

This sub-section discusses the methodology to analyze the impact of the three methods on forecast accuracy. The case study utilizes the discussed forecast accuracy metrics to compare the forecasting methods for four time periods (Year 1, Year 2, Year 3, and Year 4). However, it is important to note that year 4 is not a full annual year worth of data. Furthermore, the forecast accuracy analysis will answer four major questions:

1. Is the proposed Bayesian model valuable for forecasting all parts? (Parts are not grouped)
2. For what type of parts does the proposed Bayesian model's forecast work best? High or low demand parts?
3. Is the proposed weighting scheme for prior parameters effective?
4. When looking specifically at the LRU network, what method performs best? These results will be compared to the performance results later in this research.

The steps to answer these questions will now be discussed.

First, the case study will determine whether the Bayesian method performs superior when forecasting all parts. The forecast accuracy metrics will compare (1) all parts with and without demand and (2) all parts with demand. These results will be the basis of whether the Bayesian model performs superior to the non-Bayesian models in terms of forecast accuracy.

Second, the case study will evaluate what type of parts the Bayesian method performs superior on: high versus low demand parts. Parts will be grouped by number of demands. A high demand part is defined as a part that has greater than 3 demands in a given period, and a low demand part is defined as a part that has less than or equal to 3 demands in a given period. Bayes' has been used in previous literature when demand data is limited, which is commonly the case for low demand parts. This portion will investigate whether the proposed Bayesian model performs superior for parts with limited data (like in previous literature) by comparing the forecast accuracy of low demand parts to high demand parts.

Third, the case study will analyze whether the proposed weighting scheme for prior parameters is effective. The weighting scheme was based on the manager's confidence level in engineering estimates. Engineering estimates come from a variety of sources, and the manager's specified confidence levels based on these sources is shown in Table 3.1. This section will (1) group parts based on these sources and (2) compare the forecast accuracy of each method by source. The groups are:

1. Standard parts ($w = 2$)
2. Commercial common and consistent with commercial utilization parts ($w = 1.95$)
3. Commercial common and not consistent with commercial utilization parts ($w = 1.75$)
4. Off the shelf parts ($w = 1.95$)
5. Program unique parts ($w = 1.25$)
6. Unknown parts ($w = 1.25$)

where the weight (w) for prior parameters is specified in the parenthesis. A higher weight implies the manager is more confident in the engineering estimate, and the engineering estimate (prior) will remain important in Bayes' rule for a longer period.

Fourth, the case study will focus on the LRU network and determine what method performs best for this network. Most literature use forecast accuracy or inventory optimization results to compare forecasting methods (Bacchetti and Saccani, 2012). However, these metrics do not reflect fill rate performance, and hence, performance evaluation is recommended. Thus, this thesis seeks to compare forecast accuracy results with fill rate results, and the simulation model used to evaluate the methods' impact on fill rate can only handle LRU parts.

4.4 Impact on Inventory

The second portion of the case study determines the value of using Bayes' rule on inventory. The forecasting methods' impact on inventory stock levels and costs are compared. The strategic inventory optimization model used to optimize inventory is discussed first. Then, the methodology to analyze the methods' impact on inventory is presented.

Additionally, the remainder of the case study focuses on the Line Replaceable Unit (LRU) network. The majority of inventory is dedicated to this network, and the simulation used to compute performance is designed for these parts. By focusing on the LRU network, the impact on inventory can easily be compared to performance.

4.4.1 Strategic Inventory Optimization

The goal of this thesis is to not change how inventory is optimized. Rather the goal is to improve demand forecasting so that inventory is more accurately optimized. Hence, the model discussed in this section will be used to evaluate stock levels for all forecasting methods.

Because inventory optimization is not the focus of this thesis, the remainder of this section will briefly discuss the inventory optimization approach and refer the reader to a publication that thoroughly discusses the approach.

The commercial inventory optimization model used in this study is SPO Strategy. SPO Strategy was developed by MCA Solutions of Philadelphia, PA, and is now owned by PTC of Needham, MA. SPO Strategy is currently used at the company to optimize inventory, and this program will be used to compute inventory stock levels for the forecasting methods of interest. The inputs for this model are mean time between demands, repair time to fix a broken part, condemnation rate, procurement lead time to buy a new part, and unit price.

The objective of the strategic inventory optimization model is to determine the right mix of parts at bases and depots by minimizing back orders and inventory investment costs subject to a system's target fill rate goal. The demand in the model follows a Poisson process, and fill rate is easily computed using this assumption. In order to determine stock levels based on the objectives, the marginal analysis algorithm found in Sherbrooke (2004) is used. Sherbrooke (2004) proves that this algorithm produces an optimal backorder-versus-cost curve. He also notes that this approach has been utilized for many years since introduced in Gross (1956). This approach uses one value in each step of the algorithm to determine whether the next part should be stocked. This value is equal to the increase in overall effectiveness achieved when another unit of an item is bought. In other words, what is the next best part to select that results in the greatest "bang for the buck?" This incremental approach is how the method received its' name. Additionally, it is important to note that some parts are stocked based on their specific fill rate goal, and the remainder of the parts are optimized using the discussed algorithm.

It is important to note that the inventory model used does not necessarily mean aircraft availability is optimized. In other words, a higher overall fill rate does not necessarily optimize aircraft availability. It is the right mix of parts that achieves increased aircraft availability. However, the program of interest only manages a sub-section of parts, so optimizing aircraft availability is challenging. If the program managed all parts, multi-indenture analysis that compares the tradeoffs within various parts could be incorporated in the model. This would allow for inventory to be optimized using aircraft availability.

In conclusion, SPO Strategy is used to compute inventory stock levels. This model is a strategic optimization model where inventory levels are optimized using a marginal analysis algorithm. This method minimizes back orders and inventory costs subject to achieving a specified target fill rate goal, and fill rate is calculated using the assumption that demand is Poisson. This section notes that achieving a target fill rate goal does not necessarily result in optimized aircraft availability. However, this program does not manage all parts, so it would be difficult to optimize inventory in order to maximize aircraft availability. Hence, only fill rate is analyzed.

4.4.2 Impact on Inventory Methodology

The methodology used to determine the value of the proposed Bayesian model based on inventory is briefly discussed. The case study utilizes SPO Strategy to compute inventory stock levels and stock-holding costs for the three forecasting methods: Engineering Estimate, Current Method, and Bayesian Method. When optimizing inventory to achieve an 80% fill rate goal, mean time between demands (MTBD) is the only value varying between the methods each year. However, a difference in MTBDs does not necessarily mean a difference in target stock levels (TSLs). Thus, this section will compare the forecasting methods' MTBDs with TSLs first.

Then, the methods' impact on cost of optimal inventory is presented. This will later be compared to fill rate, the performance metric of interest.

4.5 Impact on Overall Supply Chain Performance

The third section of the case study compares the method's impact on supply chain performance. A supply chain simulation is used to understand the impact of performance over time when using Engineering Estimates, Current Method, or Bayesian Method. The simulation used is discussed first. With every simulation model, it is important to validate whether the model is accurately representing the desired environment. Thus, the steps taken to validate the simulation model used in the case study will be described. After the model is validated, a fill rate analysis can be performed. The methodology to evaluate the method's impact on supply chain performance is explained last.

4.5.1 Discrete-Event Warehouse Simulation

The Discrete-Event Warehouse Simulation employed in this research is described in Saylor and Dailey (2010). This simulation is written in the ExtendSim language and is also referred to as warehouse simulation or supply chain simulation. For the case study, the model was modified to (1) accept a user defined empirical demand distribution and (2) allow stock levels and empirical demand distributions to update periodically. This allows the study to run a multi-year simulation with varying stock levels and demand distributions by period. The concept of a data driven warehouse simulation is described in Diamond et al. (2010). The reader is referred to the discussed articles for additional details.

The simulation model has two components: supply chain and warehouse. The supply chain portion manages failing parts based on the part's mean time between demands (MTBD) or an empirical demand distribution. Failed parts are either repaired in the repair turnaround time

or condemned with a replacement part ordered a procurement lead time away. The warehouse component supplies replacement parts from warehouses and then, orders up to the target stock level when the inventory position drops to or below the re-order point. The stock level and re-order point for each part are determined by SPO Strategy, the inventory optimization model described in 4.4.1.

4.5.2 Model Validation

The purpose of this section is to validate that stock levels generated by the inventory optimization model, SPO Strategy, achieves the desired fill rate goal in the simulation model. The process to validate the simulation model is illustrated in Figure 4-3.

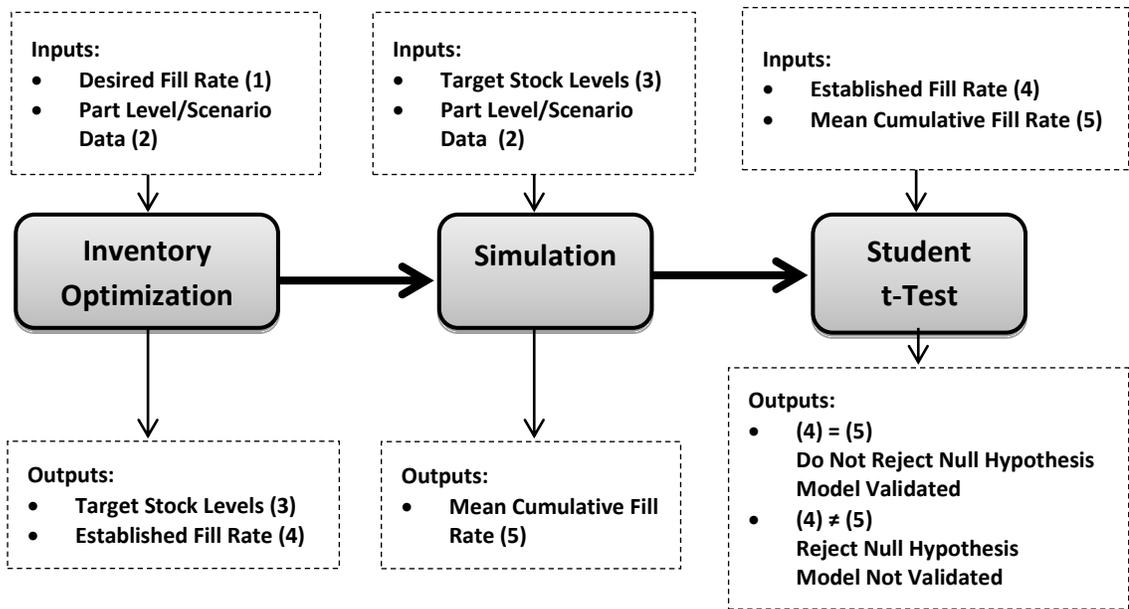


Figure 4-3 Simulation Model Validation Methodology

The inventory optimizer computes target stock levels and an established fill rate goal to achieve a minimum of the desired fill rate goal (80% in this study). These target stock levels are inputs for the simulation, and the part level data and scenario used in the inventory optimizer (such as SPO MTBD, unit price, repair turnaround time, condemnation rate, and procurement lead time)

remain the same in the simulation. It is important to note that MTBDs will be used to represent actual demand instead of an empirical distribution. A 10 year steady state simulation is conducted. The model is steady state because the scenario per month (operating hours per month and number of aircraft programs per month) remains constant. Then, if the simulation fill rate is equivalent to the established fill rate goal from the inventory optimizer, the simulation is validated. The null hypothesis is that the mean fill rate of ten warehouse simulations equals the established fill rate goal from the inventory optimizer and will be tested using a Student's t-Test with a 95% confidence level.

The assumptions for the model validation are:

1. Demand follows a Poisson distribution (variance = mean) for SPO Strategy and the demand generator in the supply chain simulation.
2. The Line Replaceable Units (LRU) network contains 239 parts.
3. SPO Strategy contains stock levels for 27 parts with demand but no price, which were individually set to achieve 80% fill rate.
4. SPO Strategy contains 30 parts with an override assigned to ensure a minimum 90% fill rate.
5. SPO Strategy was configured to ensure that each part is assigned a minimum stock level that is equal to the total forecast over the effective lead time for the part. The effective lead time is a weighted average of the repair time * percent of parts repaired + the procurement lead time * percent of parts condemned.
6. The MTBDs used in SPO Strategy were 100% causal forecasts (average operating hours per demand). The causal forecasts are based on historical demand from Year 1 to Year 4.

4.5.3 Impact on Supply Chain Performance Methodology

The Discrete-Event Warehouse Simulation model will be used to evaluate the impact of the forecasting approaches on supply chain performance. The supply chain performance metric of interest is fill rate (also known as part availability). The forecasted stock levels and observed demand are inputs in the simulation model, and these inputs are compared using the jackknife approach, which will be discussed first. After understanding these inputs, the steps and assumptions for evaluating fill rate performance are presented. Fill rates of the different forecasting methods will be compared to (1) the method's impact on inventory investment costs and (2) demand forecast accuracy results. Next, the steps followed to accomplish these tasks are outlined.

4.5.3.1 Jackknifed Approach

Jackknifed datasets are frequently used to compare different forecasting methods and will be used in this study. The jackknifing technique compares observed demand in the current period of interest with forecasts based on data from the previous period. It is important to emphasize that the observed demand in the current period is not cumulative. It is the demand seen only in that period. This approach was first mentioned in Tukey (1958) and further described in Abdi and Williams (2010). The purpose of the jackknife technique is to eliminate bias. For example, suppose a program started in Year 1. The forecast for Year 3 is based on historical demand from Year 1 to Year 2. If the forecast of Year 3 is compared with observed demand from Years 1-3, there would be a bias because the forecast of Year 3 was based on observed data from Years 1-2. The jackknife approach would compare the forecast of Year 3 with observed data only in Year 3. This would eliminate the bias. This approach will be used in the forecast accuracy portion, and the implementation in forecast accuracy is straightforward. However, in terms of the simulation, the implementation needs further explanation. Demand is

split in two components to create the jackknifed forecast. These components are (1) forecasted demand in the period of interest and (2) historical demand in the period of interest. The data used for these components in the simulation will now be explained.

The forecasted demand in the period of interest is represented by the target stock levels evaluated in SPO Strategy. For example, suppose the forecasting approach is based on historical data. If the period of interest is Year 3, the optimized target stock levels for Year 3 are based on the mean time between demand (MTBD) due to observed demand in Year 1 and Year 2.

The historical demand in the period of interest is represented by an empirical distribution. If the example from above is used, the period of interest is Year 3. The empirical distribution would be based on actual demand in Year 3. An empirical distribution is used instead of actual demand because it allows for the ability to run multiple scenarios in order to calculate a confidence interval and run a steady-state scenario for future periods. The empirical distribution is created for each part by a three step process:

1. Sum the historical demands per month for each part number, including months with zero demand.
2. Create a cumulative histogram with unequal bins for each part number where each bin corresponds to the actual monthly demand. The resulting histogram shows the cumulative probability of having 0 to n demands. This histogram is shown in Figure 4-4 where the x-axis is demands per month and the y-axis is the cumulative probability. Although histograms usually have equal bin sizes, for creating an empirical demand distribution, this histogram uses bins that are equal to the actual monthly demands that occur. The number of bins is equal to the number of unique values of demands per month (plus one if there are months with zero demand).

3. Each month, for each part, the demand generating function in the supply chain simulation draws a random number between 0 and 100 percent and looks up the cumulative probability in the cumulative histogram to determine the corresponding demands per month. In this way, the demand generating function randomly determines monthly demands per part following the same empirical demand distribution as the actual data.
4. Repeat steps for each period.

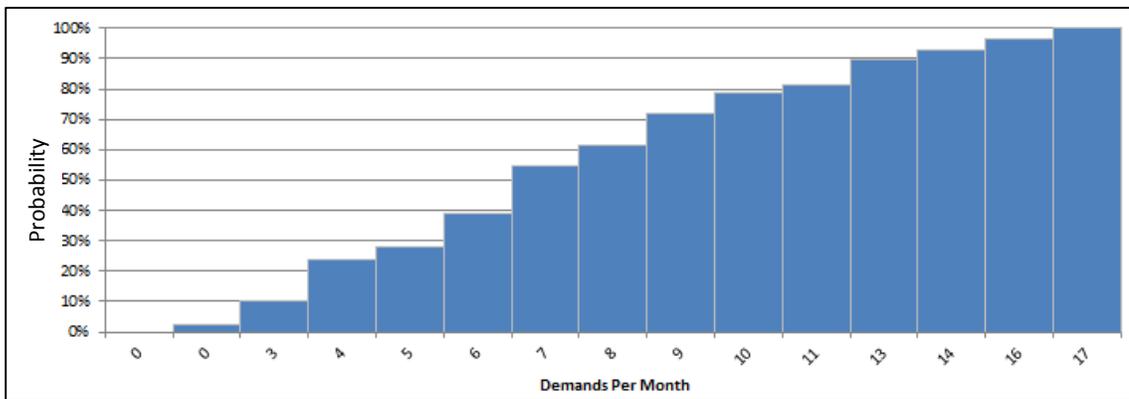


Figure 4-4 Cumulative Histogram for Demands per Month

4.5.3.2 Simulation Methodology and Assumptions

The supply chain simulation is used to compare the method's impact on fill rate over time. The steps and assumptions used in the simulation are described first. Then, the analysis on fill rate presented in the case study is explained.

The simulation model allows the stock levels and empirical demand distribution to change periodically (annually in this study), so only one simulation model for each method is needed. The simulation process involves two steps for each method:

1. Input data into simulation: target stock levels based on the desired forecasting method, empirical distribution based on observed demand, scenarios (number of aircraft and operating hours each month), part's unit price, repair turnaround time, condemnation rate, and procurement lead time.
2. Execute multiple runs of the simulation model. 30 runs with an 80% confidence level will be used in this study.

Additionally, there are a few assumptions used in the simulation. These are:

1. Observed demand follows an empirical distribution. For more detail, refer to 4.5.2.
2. Empirical demand distribution will vary each year from Year 1 to Year 4. Years 1-3 will use the year's annual demand data to create the distribution. As Year 4 only contains three months of historical data, the period from 01/year 1 – 03/year 4 is used for this year. For Year 5 and beyond, the period from 01/year 1 – 03/year 4 is used in order to ensure that all parts that have experienced actual historical demand are represented in the empirical demand distribution.
3. At the start of the program, engineering estimates are used to forecast demand. Thus, the inventory levels by recommended by the engineering estimates (taking procurement lead time of each part into account) are on hand at the start of the program in Year 1.
4. Each subsequent year (Year 2, Year 3, Year 4), stock levels will change, and if additional spares are required, then spares are purchased and arrives a procurement lead time away. If there is excess of a spare, excess stock will be used.
5. Part level data and scenario (number of aircraft and total operating hours) for simulation are equivalent to data in SPO Strategy, inventory optimizer.

The supply chain simulation is used to understand the method's impact on supply chain performance over time, measured by the fill rate. The above steps and assumptions are used to compute fill rate, and once fill rate is computed, it is compared to the case studies previous results. The overall methodology can be summarized as:

1) Fill rate will be compared to inventory investment costs. A lower inventory investment cost at a specified fill rate goal does not necessarily mean an approach is superior, and as such, this comparison is important.

2) Fill rate is compared to the demand forecast accuracy results. Forecast accuracy is frequently used in comparative studies, and this analysis investigates whether forecast accuracy is a good indicator of what forecast method is "best."

Chapter 5 – Case Study Results

5.1 Overview of Case Study

The purpose of the case study is to validate the use of Bayes' rule when estimating the demand of spare parts. The case study is performed on an aircraft program with 3 ¼ years of data. Three forecasting methods are compared in the case study: Engineering Estimates, Current Method, and Bayesian Method. These methods will be compared in three sections: the forecasting method's impact on (1) forecast accuracy, (2) inventory, and (3) performance, respectively. Additionally, the analysis on forecast accuracy investigates parts in all networks (LRU, Consumable, MSC) while analysis on inventory and performance only consider the LRU network. The results of the case study will now be discussed.

5.2 Impact on Forecast Accuracy

The first section of the case study determines the value of using the proposed Bayesian method on forecast accuracy. The forecast accuracy for all parts will be compared first. Then, this thesis will explore what type of parts Bayes' forecasts superior for, low versus high demand parts. After, the proposed weighting scheme for prior parameters will be evaluated using the accuracy metrics, and last, the forecast accuracy of the Line Replaceable Unit (LRU) network will be analyzed. The demand forecast accuracy metrics used are Mean Absolute Deviation (MAD), the ratio of MAD to the mean (MAD/mean), and Mean Absolute Percentage Error (MAPE). These metrics will be compared with respect to annual time periods because the majority of parts experience low demand.

5.2.1 Forecast Accuracy

The case study will investigate whether the Bayesian model performs superior to the other methods when forecasting demand for all parts. Table 5.1, Table 5.2, Table 5.3, and Table 5.4 illustrate the forecast accuracy in Year 1, Year 2, Year 3, and Year 4, respectively, where “n” indicates the total number of parts. Also, for all forecast accuracy metrics, a lower value is better meaning a more accurate forecast.

Table 5.1 Forecast Accuracy in Year 1

	Year 1					
	All Parts (n = 439)			Parts with Demand (n = 36)		
	EE	Current	Bayes'	EE	Current	Bayes'
MAD	0.93	0.93	0.93	0.93	0.93	0.93
MAD/Mean	2.12	2.12	2.12	0.17	0.17	0.17
MAPE				98.7%	98.7%	98.7%

Table 5.2 Forecast Accuracy in Year 2

	Year 2					
	All Parts (n = 439)			Parts with Demand (n = 57)		
	EE	Current	Bayes'	EE	Current	Bayes'
MAD	1.11	1.11	1.17	4.47	4.47	4.87
MAD/Mean	1.29	1.29	1.35	0.67	0.67	0.73
MAPE				71.8%	71.8%	73.4%

Table 5.3 Forecast Accuracy in Year 3

	Year 3					
	All Parts (n = 439)			Parts with Demand (n = 64)		
	EE	Current	Bayes'	EE	Current	Bayes'
MAD	1.38	1.63	1.23	4.84	5.28	4.61
MAD/Mean	1.33	1.57	1.19	0.68	0.74	0.65
MAPE				85.5%	91.3%	81.1%

Table 5.4 Forecast Accuracy in Year 4

	Year 4					
	All Parts (n = 439)			Parts with Demand (n = 31)		
	EE	Current	Bayes'	EE	Current	Bayes'
MAD	0.36	0.36	0.26	2.62	2.45	1.96
MAD/Mean	1.61	1.64	1.18	0.84	0.78	0.63
MAPE				87.7%	102.17%	85.4%

In Year 1, all metrics perform the same because there was no data for Bayes' to learn from yet. Also, the Current Method uses the engineering estimates for the first two years of the program's life. In Year 2, the Engineering Estimate and Current Method perform superior to Bayes' in every forecast accuracy metric. In Year 3 and Year 4, Bayes' performs superior to both methods in every metric. Additionally, the forecast accuracy results on the Current Method are interesting. In Year 3, the forecast accuracy metrics conclude that the Current Method performs worst in Year 3. In Year 4, the Current Method performs worst for all parts (with and without demand). However, the Current Method performs in the middle for all parts with demand in Year 3. These results suggest that the Current Method does not perform well for parts with zero demands. However, the Bayesian method performs superior in all years except Year 2.

So, why does the Bayes' approach perform the worst in Year 2? When investigating the forecast accuracy data, it turns out that one part is driving the metric's results. The engineering estimate for this part is 76 operating hours per failure. However, in Year 1, the part failed much less than predicted by the Engineering Estimate (360 operating hours/failure). As a result, Bayes' estimate for Year 2 is 345 operating hours/failure, and in Year 2, the actual operating hours per failure (68 operating hours/failure) was much closer to the Engineering Estimate. Thus, Bayes' performs poorly. This investigation emphasizes some key conclusions. First, the forecast accuracy metrics can be greatly affected by one outlier. The metrics draw attention to

large absolute errors, which can be exaggerated by one high demand part when most parts are low demand. Hence, when comparing metrics with a majority of expensive low demand parts, it is important to be cognizant of this. Second, when initial data varies greatly from long-term experience, making decisions based on the initial data can lead to poor performance. The Current Method tries to adjust for this by waiting two years to utilize observed data, and Bayes' tries to adjust for this by including prior knowledge in the estimate. However, observed data can quickly diminish the importance of prior knowledge in Bayes' rule, so developing an appropriate weighting scheme for prior parameters in the proposed Bayesian model is very important. Table 5.5 illustrates the forecast accuracy results if this part was weighted properly and is equivalent to the Engineering Estimate. These results show that Bayes' performs best on all metrics besides MAPE. Additionally, although this part greatly affects the metrics, it is in the set of parts to analyze. Thus, it will still be included in the data set for the remainder of this research.

Table 5.5 Forecast Accuracy if Outlier Part was Weighted Appropriately in Year 2

	Year 2					
	All Parts (n = 439)			Parts with Demand (n = 57)		
	EE	Current	Bayes'	EE	Current	Bayes'
MAD	1.11	1.11	1.17	4.47	4.47	3.85
MAD/Mean	1.29	1.29	1.35	0.67	0.67	0.58
MAPE				71.8%	71.8%	72.2%

5.2.2 Low versus High Demand Parts

This section will explore how well Bayes' performs on high demand versus low demand parts. Bayes' has been used in literature when data is limited, which is the case when parts have low demand. Thus, it is important to analyze whether Bayes' is truly useful for parts with limited data. This research defines a low demand part as less than or equal to three demands in a

period and a high demand part as greater than three demands in a period. The period is defined as a year. Table 5.6 summarizes the forecast accuracy results on what forecasting method performs best for low and high demand parts. Table 5.7 shows the forecast accuracy results for all low demand parts, and Table 5.8 shows the forecast accuracy results for low demand parts with demand. Table 5.9 shows the forecast accuracy results for high demand parts, and Table 5.10 shows the forecast accuracy results for high demand parts with demand. A lower value is better for all forecast accuracy metrics. Also, Year 1 will no longer be explored because the forecasts for all three methods are equivalent.

Table 5.6 Forecast Accuracy for Low Demand versus High Demand Parts

		Year 2	Year 3	Year 4
All Parts	Demand ≤3	Bayes' (n=422)	Bayes' (n=420)	Bayes' (n=435)
	Demand >3	EE/Current (n=17)	Bayes' (n=19)	Bayes' (n=4)
Parts with Demand	Demand ≤3	Bayes' (n=40)	Bayes' (n=45)	Bayes' (n=27)
	Demand >3	EE/Current (n=17)	Bayes' (n=19)	Bayes' (n=4)

Table 5.7 Forecast Accuracy for All Low Demand Parts

	All Low Demand Parts								
	Year 2			Year 3			Year 4		
	EE	SPO	Bayes'	EE	SPO	Bayes'	EE	SPO	Bayes'
MAD	0.66	0.66	0.63	0.84	1.05	0.72	0.27	0.28	0.21
MAD/Mean	4.39	4.39	4.25	5.27	6.57	4.53	2.69	2.86	2.13
RMSE	0.81	0.81	0.80	0.92	1.02	0.85	0.52	0.53	0.46

Table 5.8 Forecast Accuracy for Low Demand Parts with Demand

	Low Demand Parts with Demand								
	Year 2			Year 3			Year 4		
	EE	SPO	Bayes'	EE	SPO	Bayes'	EE	SPO	Bayes'
MAD	1.05	1.05	1.01	1.30	1.40	1.26	1.51	1.46	1.40
MAD/Mean	0.67	0.67	0.64	0.87	0.94	0.85	0.95	0.92	0.88
RMSE	1.03	1.03	1.01	1.14	1.18	1.12	1.23	1.21	1.18
MAPE	0.67	0.67	0.65	0.89	0.95	0.85	0.89	1.03	0.94

Table 5.9 Forecast Accuracy for all High Demand Parts

	All High Demand Parts								
	Year 2			Year 3			Year 4		
	EE	SPO	Bayes'	EE	SPO	Bayes'	EE	SPO	Bayes'
MAD	12.50	12.50	13.94	13.22	14.48	12.52	10.10	9.14	5.78
MAD/Mean	0.67	0.67	0.75	0.65	0.71	0.61	0.75	0.68	0.43
RMSE	3.54	3.54	3.73	3.64	3.81	3.54	3.18	3.02	2.41

Table 5.10 Forecast Accuracy for High Demand Parts with Demand

	High Demand Parts with Demand								
	Year 2			Year 3			Year 4		
	EE	SPO	Bayes'	EE	SPO	Bayes'	EE	SPO	Bayes'
MAD	12.50	12.50	13.94	13.22	14.48	12.52	10.10	9.14	5.78
MAD/Mean	0.67	0.67	0.75	0.65	0.71	0.61	0.75	0.68	0.43
RMSE	3.54	3.54	3.73	3.64	3.81	3.54	3.18	3.02	2.41
MAPE	0.82	0.82	0.93	0.78	0.82	0.72	0.79	0.96	0.56

Table 5.6 shows that the majority of parts are low demand. When looking at low demand parts (demands ≤ 3), Bayes' performs superior to the other methods every year. This validates Bayes' works well for parts that have limited data. When looking at high demand parts (demands > 3), the Engineering Estimate and Current Method perform superior to Bayes' in Year 2. However, Bayes' performs superior for high demand parts in Year 3 and Year 4. Thus, this analysis not only validates the use of Bayes' rule for parts with limited data, but it supports the use of Bayes' for high demand parts the majority of the time.

5.2.3 Appropriate Weighting Scheme

This research uses forecast accuracy metrics to determine whether the weighting scheme for prior parameters is appropriate. Parts are grouped by source of the engineering estimate in this study. Table 5.11 summarizes the forecast accuracy results on what forecasting method performs best for all parts with and without demand. The column on the far left illustrates the confidence levels for each group where a higher number means more confident in the engineering estimate, and the second column from the left shows the part groups based on the source of engineering estimates. Table 5.12 presents the forecast accuracy values. A lower value is better for all forecast accuracy metrics. Also, if a value equals “N/A” that means the metric was unable to compute the value because zero observed demands exists.

Table 5.11 Forecast Accuracy Summary for Source of Engineering Estimate

		Year 2	Year 3	Year 4
100%	Standard (n=3)	Same	Bayes'	Bayes'
95%	Commerical Common - Consistent with Commercial Utilization (n=42)	EE/ Current	Bayes'	Bayes'
75%	Commerical Common - Not Consistent with Commercial Utilization (n=4)	Same	Same	Same
95%	Off the Shelf (n=37)	EE/ Current	EE	EE
25%	Unique (n=352)	Bayes'	Bayes'	Bayes'
25%	Unknown (n=1)	Same	Same	Same

Table 5.12 Forecast Accuracy Detailed for Source of Engineering Estimate

		Standard			Commerical Common - Consistent with Commercial Utilization			Commerical Common - Not Consistent with Commercial Utilization			Off the Shelf			Unique			Unknown		
		EE	SPO	Bayes'	EE	SPO	Bayes'	EE	SPO	Bayes'	EE	SPO	Bayes'	EE	SPO	Bayes'	EE	SPO	Bayes'
Year 2	Mad	0.46	0.46	0.46	3.72	3.72	4.21	0.25	0.25	0.25	0.84	0.84	1.09	0.85	0.85	0.84	1E-05	1E-05	1E-05
	Mad/Mean	0.69	0.69	0.69	0.71	0.71	0.80	N/A	N/A	N/A	1.36	1.36	1.75	2.22	2.22	2.19	N/A	N/A	N/A
	RMSE	0.68	0.68	0.68	1.93	1.93	2.05	0.50	0.50	0.50	0.92	0.92	1.04	0.92	0.92	0.92	3E-03	3E-03	3E-03
Year 3	Mad	1.25	1.25	0.89	4.17	5.09	4.02	0.29	0.29	0.29	0.94	2.80	1.49	1.11	1.12	0.89	1E-05	1E-05	1E-05
	Mad/Mean	N/A	N/A	N/A	0.67	0.81	0.64	N/A	N/A	N/A	6.93	20.73	11.01	2.08	2.10	1.68	N/A	N/A	N/A
	RMSE	1.12	1.12	0.94	2.04	2.26	2.00	0.54	0.54	0.54	0.97	1.67	1.22	1.05	1.06	0.94	4E-03	4E-03	4E-03
Year 4	Mad	0.31	0.31	0.18	1.27	1.10	0.71	0.07	0.07	0.07	0.26	0.46	0.40	0.26	0.27	0.20	3E-06	3E-06	3E-06
	Mad/Mean	N/A	N/A	N/A	0.98	0.85	0.55	N/A	N/A	N/A	1.95	3.37	2.97	2.43	2.51	1.82	N/A	N/A	N/A
	RMSE	0.55	0.55	0.43	1.13	1.05	0.84	0.27	0.27	0.27	0.51	0.67	0.63	0.51	0.52	0.44	2E-03	2E-03	2E-03

When evaluating “Standard” and “Unique” parts, Bayes’ performs superior. Thus, the weighting scheme seems appropriate for these parts. It is important to note that “Unique” parts are the majority of parts, and managers have very low confidence in the engineering estimates. Hence, Bayes’ works well for these parts where managers have a low SPO confidence in the engineering estimate. Next, when analyzing “Commercial Common – Consistent with Commercial Utilization” parts, Bayes’ performs superior in all years except Year 2. Hence, a higher weight on the engineering estimate might appropriate for these parts. Further, Bayes’ performs the same as the other methods for “Unknown” and “Commercial Common – Consistent with Commercial Utilization” parts because the prediction of demand for all methods is equivalent. Last, Bayes’ performs inferior to the Engineering Estimate each year for “Off the Shelf” parts. Hence, higher weight on “Off the Shelf” prior parameters is very apparent.

5.2.4 LRU Forecast Accuracy

The case study will now focus only on the LRU network and determine what method performs superior for this network. Table 5.13 summarizes which method performs best (1), second best (2), and worst (3), and shows the forecast accuracy results. Table 5.14 illustrates the forecast accuracy values, and a lower value is better for all forecast accuracy metrics.

Table 5.13 Summary of Forecast Accuracy of LRU Parts

Year 2	Year 3	Year 4
1) EE/Current 3) Bayes'	1) Bayes' 2) EE 3) Current	1) Bayes' 2) EE 3) Current

Table 5.14 Forecast Accuracy of LRU Parts

	Year 2			Year 3			Year 4		
	EE	SPO	Bayes'	EE	SPO	Bayes'	EE	SPO	Bayes'
MAD	1.29	1.29	1.41	1.55	1.99	1.48	0.44	0.45	0.33
MAD/Mean	0.94	0.94	1.02	0.96	1.23	0.92	1.22	1.26	0.92
RMSE	1.14	1.14	1.19	1.24	1.41	1.22	0.66	0.67	0.58

Table 5.13 indicates that Bayes’ performs worst in Year 2 and best in Year 3 and Year 4.

Interestingly, in Year 3 and Year 4, the Engineering Estimate performs second best and the Current Method performs worst.

5.2.5 Forecast Accuracy Conclusion

The proposed Bayesian approach (Bayes’) was found to perform superior in each year except Year 2. Further investigation into this shortcoming found that the forecast accuracy metrics were heavily influenced by one outlier part. When analyzing the outlier part, it was found that the engineering estimate for this part was quickly overwhelmed by observed data in Bayes’ rule. This emphasizes the importance of developing an appropriate prior parameter weighting scheme. Next, this section evaluates whether Bayes’ works well with high versus low demand parts and validates that Bayes’ works well for parts with limited data (or low demand parts). After, the weighting scheme for prior parameters is analyzed. From the results, the only group that seemed necessary to modify was “Off the Shelf” parts. Last, the forecast accuracy of the LRU network is evaluated. Bayes’ performs best for this network each year except Year 2. In

Year 2, the Engineering Estimate performs best. Also, interestingly enough, the Current Method performs worst in Year 3 and Year 4. This analysis would advise a forecaster to not use the Current Method.

5.3 Impact on Inventory

The second section of the case study determines the value of the proposed Bayesian model on inventory. Optimal stock levels and inventory investment costs for each annual year are found using SPO Strategy. The only varying component in the model each year is the mean time between demands (MTBD) between forecasting methods: Engineering Estimates, Current Method, and Bayesian Method. This section focuses on the LRU network, and the model is optimized to achieve an 80% fill rate goal. The method's mean time between demands will be compared to target stock levels (TSLs) first. Then, the impact of the methods on cost of optimal inventory is presented.

5.3.1 Mean Time between Demands versus Stock Levels

Mean time between demands (MTBD) is found by the forecasting methods, and these MTBDs are inputs for inventory optimization, which evaluates target stock levels (TSL). However, a difference in MTBDs does not necessarily mean there is a difference in target stock levels. Here, the forecasting's methods MTBDs with TSLs are compared each year. Table 5.15 shows the percentage of parts in each of the five possible combinations of values (MTBDs or TSLs) each year. These combinations are: "Bayes=Current=EE", "Bayes<>Current=EE", "Bayes=EE<>Current", "Bayes=Current<>EE", and "Bayes=Current<>EE." "=" means the methods are equal in value, and "<>" means the methods are not equal in value. For example, a part would be in "Bayes<>Current=EE" category if the MTBD for the Current Method and Engineering Estimate equal 500 operating hours per demand, and the MTBD for Bayes' equals

200 operating hours per demand. Also, the total percentage for all five groups each year should equal 100% if all parts are considered.

Table 5.15 Mean Time between Demands (MTDB) versus Target Stock Levels (TSL)

	MTBD				TSL			
	Year 1	Year 2	Year 3	Year 4	Year 1	Year 2	Year 3	Year 4
Bayes=Current=EE	100%	87%	74%	64%	100%	83%	84%	80%
Bayes<>Current=EE		13%	18%	25%		17%	1%	6%
Bayes=EE<>Current								3%
Bayes=Current<>EE							13%	7%
Bayes<>Current<>EE			8%	11%			1%	4%
TOTAL	100%	100%	100%	100%	100%	100%	100%	100%

However, Table 5.15 is difficult to interpret when comparing two methods. Thus, Table 5.16, Table 5.17, and Table 5.18 are presented to easily make comparisons between two methods. Table 5.16 compares the Bayesian Method (Bayes) and the Engineering Estimate (EE). Table 5.17 compares the Current Method (Current) and Engineering Estimate (EE), and Table 5.18 compares Current Method (Current) and the Bayesian Method (Bayes). These tables can be interpreted the same as Table 5.15 where “=” means the forecasting methods are equal in value and “<>” means the forecasting methods are not equal in value.

Table 5.16 Bayesian Method versus Engineering Estimate Mean Time between Demands (MTBD) and Target Stock Levels (TSLs)

	MTBD				TSL			
	Year 1	Year 2	Year 3	Year 4	Year 1	Year 2	Year 3	Year 4
Bayes = EE	100%	87%	74%	64%	100%	83%	84%	83%
Bayes <> EE	0%	13%	26%	36%	0%	17%	15%	17%
TOTAL	100%	100%	100%	100%	100%	100%	100%	100%

**Table 5.17 Current Method versus Engineering Estimate Mean Time between Demands
(MTBD) and Target Stock Levels (TSLs)**

	MTBD				TSL			
	Year 1	Year 2	Year 3	Year 4	Year 1	Year 2	Year 3	Year 4
Current = EE	100%	100%	92%	89%	100%	100%	85%	86%
Current <> EE	0%	0%	8%	11%	0%	0%	14%	14%
TOTAL	100%	100%	100%	100%	100%	100%	100%	100%

**Table 5.18 Bayesian Method versus Current Method Mean Time between Demands (MTBD)
and Target Stock Levels (TSLs)**

	MTBD				TSL			
	Year 1	Year 2	Year 3	Year 4	Year 1	Year 2	Year 3	Year 4
Bayes = Current	100%	87%	74%	64%	100%	83%	97%	87%
Bayes <> Current	0%	13%	26%	36%	0%	17%	2%	13%
TOTAL	100%	100%	100%	100%	100%	100%	100%	100%

When looking at Table 5.16 and Table 5.17, one can see that the Bayesian Method's MTBDs are moving away much faster from the Engineering Estimate than the Current Method's MTBDs. Once the Engineering Estimate is revised by Bayes' in Year 2, the target stock levels immediately vary from the Engineering Estimate's TSLs by 17%. However, as Bayes' MTBDs increasingly differ from the Engineering Estimate, Bayes' target stock levels remain steady from Year 2-Year 4 with only a 15 - 17% difference from the Engineering Estimate's TSLs. This same pattern occurs with the Current Method. Once the Current Method revises the Engineering Estimate with historical demand in Year 3, the MTBDs start to move away from the Engineering Estimate; however, once the target stock levels change in Year 3 with a 14% difference, the difference in TSLs remains steady. Thus, one can conclude that the target stock levels are most effected when the estimates are first revised with observed demand.

When comparing the Bayesian Method to the Current Method in Table 5.18, the MTBDs seem to move away from each other as the years increase. However, it is important to note that Bayes' MTBDs are actually moving towards the Current Method each year, but Bayes' holds on to engineering estimates. Thus, the MTBDs might be similar but not be exactly the same, and this is the reason the MTBDs seem to be moving away from each other in Table 5.18. The target stock levels in this figure verify this conclusion. The difference in MTBDs between the methods increases from 13% in Year 2 to 26% in Year 3, but the difference in target stock levels between the methods decreases from 17% in Year 2 to 2% Year 3.

5.3.2 Inventory Investment Costs

The cost of inventory to achieve an 80% fill rate goal is evaluated for each of the three forecasting techniques. These results are presented in Figure 5-1 where Engineering Estimates is denoted as EE, Current Method is denoted as Current, and Bayesian Method is denoted as Bayes.



Figure 5-1 Optimal Inventory Investment Costs

The aircraft program is first employed in Year 1. As a result, the only values for MTBD are from Engineering Estimates. Thus, all forecasting methods use the Engineering Estimate to optimize inventory and end with the same target stock levels. The starting inventory position is the same for all methods at \$8,582,095.

In Year 2 (the second year of operation), the Current Method uses Engineering Estimates to estimate MTBD. The business rule in the Current Method states that engineering estimates will transition to actual demand after two years of observed data. Thus, the Current Method and Engineering Estimate stay at an inventory investment of \$8,582,095. Bayes' revises the engineering estimates in light of Year 1's observed demand, and as a result, the inventory investment for Bayes' increases to \$10,760,564.

In Year 3 (the third year of operation), both the Current Method and Bayesian Method revise stock levels due to actual demand. The business rule for the Current Method states that after two years of observed data, demand will transition from engineering estimates to observed demand when there are 3 or more demands in a 12 month period. As a result, the stock holding costs for the Current Method increases to \$10,868,119. Bayes' is influenced by all parts with demand and the Engineering Estimates. As a result, the inventory investment for Bayes' Rule drops slightly to \$10,750,357. The Engineering Estimate's inventory investment remains static at \$8,582,095.

In Year 4 (the fourth year of operation), the historical demand data does not represent an entire annual year of data. The data exists from January in Year 3 through March in Year 4. However, there is still a significant change in inventory investment for the Current Method and Bayesian Method. The Current Method drops by approximately \$1,137,000 to \$9,731,252 in

inventory investment. The Bayesian Method drops by approximately \$1,662,000 to \$9,087,943 in inventory investment.

This analysis looks only at annual inventory investment costs optimized at an 80% fill rate. Based on these results, the Engineering Estimate would return the lowest inventory investment at the specified fill rate goal each year. However, these investments will be compared to the simulated fill rate in the next section.

5.4 Impact on Fill Rate Performance

The last section of the case study uses the Discrete-Event Warehouse Simulation found in Saylor and Dailey (2010) to evaluate the impact of the forecasting approaches on performance. The performance metric of interest is fill rate (also known as part availability). First, the simulation model must be validated, and this is presented first. Once validated, fill rate analysis can be achieved. Fill rates of the different forecasting methods will be compared to (1) the method's impact on inventory investment costs and (2) demand forecast accuracy results. The results are presented next, respectively.

5.4.1 Model Validation

The purpose of the model validation is to verify that the stock levels generated by the inventory optimization model, SPO Strategy, achieves the desired fill rate goal in the simulation model. Stock levels are optimized in SPO Strategy and loaded into the simulation. After conducting the simulation, a Student t-Test was used to determine whether the specified optimization fill rate equals the simulation fill rate.

SPO Strategy is configured to compute target stock levels for parts grouped in the Line Replaceable Unit (LRU) network. The objective function is to minimize inventory investment

and back orders subject to the constraint of achieving an 80% target fill rate goal. After running the optimizer, SPO Strategy achieved an 80.58% fill rate goal at the minimum inventory investment of \$10,529,384. Next, these computed stock levels are loaded in the simulation. The scenario (number of aircrafts and flight hours per aircraft) and part level data (MTBD, price, repair turnaround time, condemnation rate, and procurement lead time) remain the same in the simulation. Ten simulations were run and the mean cumulative fill rate goal for the simulation is 80.45%. Table 5.19 shows the cumulative fill rates for the ten simulation run.

Table 5.19 Simulated Cumulative Fill Rate over Ten years

Simulation Run	Cumulative Fill Rate
1	81.4%
2	82.0%
3	80.6%
4	79.6%
5	78.6%
6	79.2%
7	81.5%
8	80.3%
9	79.8%
10	81.5%

A Student t-Test is used to test whether the simulation's mean cumulative fill rate goal (80.45%) is equivalent to the desired optimization fill rate goal (80.58%). The null hypothesis is that the mean of the ten warehouse simulations equals the fill rate goal established in the inventory optimization model. The test was run at the 5% significance level, meaning that the null hypothesis is only rejected when it is true 5% of the time; this is termed a Type I error in statistics. At a 95% confidence level, there is no reason to reject the null hypothesis. Thus, the warehouse simulation fill rate (80.45%) achieves the same fill rate as SPO Strategy (80.58%), and the simulation model is validated. The inputs and outputs for the Student t-Test are illustrated in Table 5.20 and Table 5.21, respectively.

Table 5.20 Student's t-Test Inputs for the Null Hypothesis that Simulated Fill Rate is Equal to the Optimization Fill Rate of 80.58%

Input Name	Input Value
Null Hypothesis (H_0)	$\mu = .0858$
Alternative Hypothesis (H_1)	$\mu \neq .0858$
Level of Significance (α)	0.05
Sample Size (n)	10
Sample Mean (x)	0.804467
Sample Standard Deviation (s)	0.011428

Table 5.21 Student's t-Test Outputs for the Null Hypothesis that Simulated Fill Rate is Equal to the Optimization Fill Rate of 80.58%

Output Name	Output Value
Test statistic	-0.3689
P-value	0.721
Reject Null Hypothesis?	P-value (.721) > Level of Significance (α) DO NOT Reject Null Hypothesis ($\mu = .0858$)

5.4.2 Fill Rate Performance versus Inventory Investment

Supply chain simulation is used to understand the forecasting method's impact on fill rate over time. The fill rates over time for each method are given in Figure 5-2 where Engineering Estimates are shown in red, Current Method is shown in orange, and Bayesian Method is shown in green. Figure 5-2 indicates that overall the Engineering Estimates approach performs significantly inferior to the other methods. Thus, the Current and Bayesian Method's fill rates are further compared to inventory investment costs (stock-holding costs) in Figure 5-3. Figure 5-3 highlights the method's starting and ending fill rate performance from year 2 to year 4.

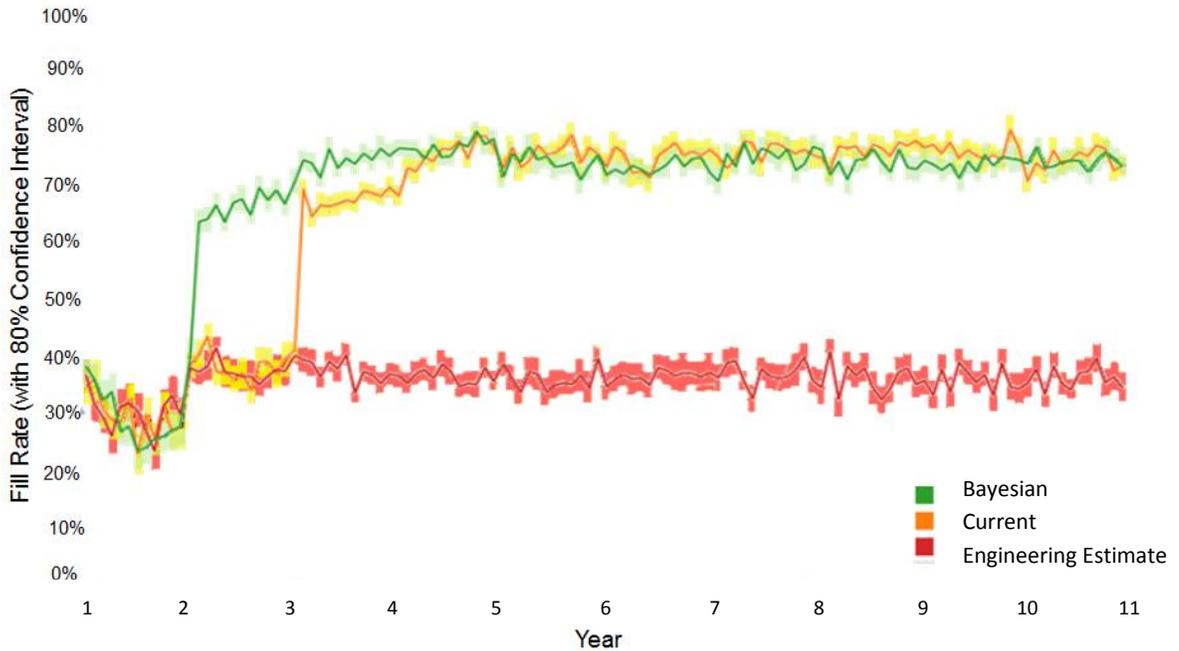


Figure 5-2 Chart of fill rate over time, with an 80% confidence interval, for Bayesian Method (Green), Current Method (Orange) and Engineering Estimates (Red). Higher fill rate is better.

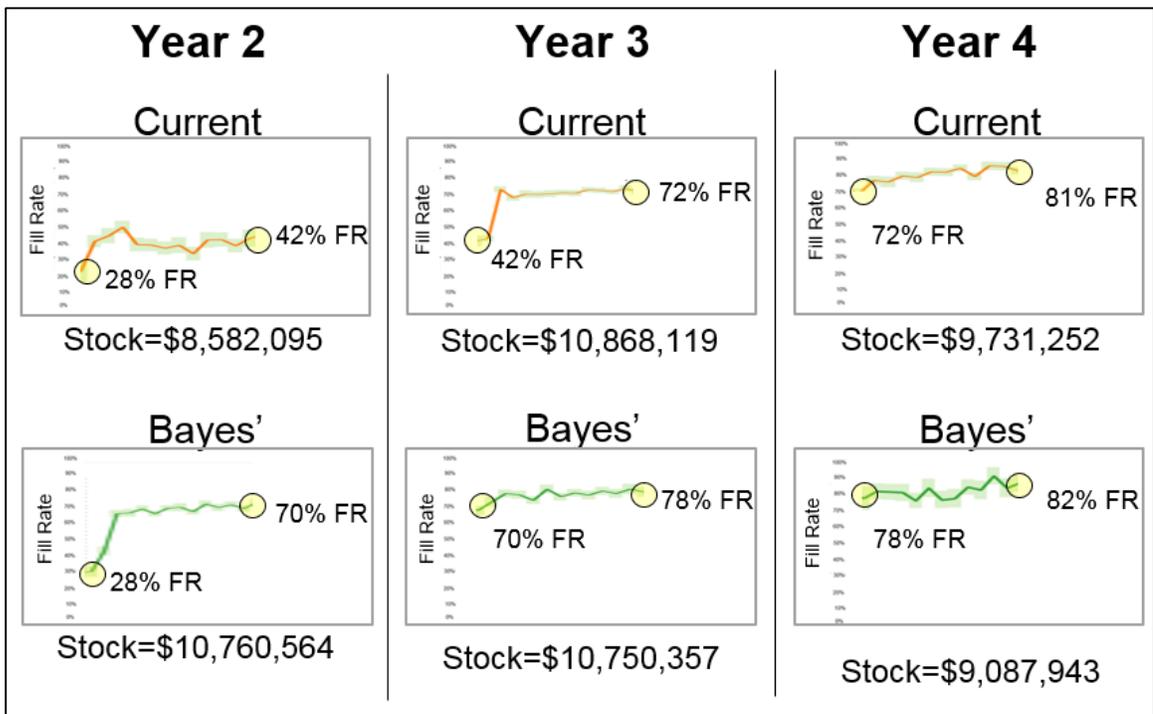


Figure 5-3 Fill Rate Performance Comparison of Current Approach versus Bayesian Approach

As seen in Figure 5-2, all forecasting methods rely on the engineering estimate for the first year of operation, Year 1. The inventory investment of \$8,582,095 is supported by a fill rate around 30%.

In the second year of operation, Year 2, the Current Method relies on engineering estimates to set stock levels. This is because the business rule requires two years of historical data before observed demand is used. Engineering Estimate and the Current Method remain constant in terms of inventory investment cost (\$8,582,095) and slightly increase in fill rate performance (starting the year at 28% and ending the year at 42%). The Bayesian Method revises the engineering estimate in light of actual demand, and the inventory investment increases to \$10,750,357; however, fill rate also significantly increases (starting the year at 28% and ending the year at 70%). This explains the increase of inventory investment and emphasizes the need to budget for increased inventory investment on new aircraft programs as estimates of demand are refined over time.

In the third year of operation, Year 3, the Current Method uses historical data when parts have 3 or more demands in a 12 month period. Due to this, the Current Method's fill rate significantly increases (starting the year at 42% and ending the year at 72%) with a required inventory investment of \$10,868,119. The gap between the Current Method and Bayes' significantly reduces. However, Bayes' performs better in terms of fill rate (starting the year at 70% and ending the year at 78%) with a slightly lower inventory cost of \$10,750,357.

In Year 4 and beyond, stock levels based on 03/year 1 through 03/year 4 data were evaluated against an empirical demand distribution from this same time period. It can be argued that this is an unfair assessment, as the same dataset used to train the demand forecast is used to evaluate it. Nonetheless, in year 4 Bayes' (starting the year at 78% and ending the

year at 82%) performs slightly better than the Current Method (starting the year at 72% and ending the year at 81%), and Bayes' has around a \$700K lower inventory cost. In the middle of year 4, the Current Method and Bayesian Method start to yield similar performance, and this trend continues for the rest of the simulation.

Three key conclusions come from this analysis. First, the Current Method and Bayesian Method both outperform Engineering Estimates in supporting aircrafts. Over time, both the Current Method and Bayesian Method result in similar performance. However, Bayes' significantly outperforms the Current Method in the first few years of the program because it incorporates the experience of early demands. Thus, learning from demand immediately is important. Second, when engineering estimates were revised by actual demand, fill rate along with inventory investment significantly increased. This was seen for Bayes' in Year 2 and for Current Method in Year 3. Hence, management should absolutely budget for increased inventory investment costs on new aircraft programs as estimates of demand are refined over time. Third, this case study shows the importance of evaluating true fill rate performance in comparative studies. If this analysis was not performed, the Engineering Estimate would have seemed superior when looking at inventory investment costs. The investment was lower each year at the optimized fill rate goal. However, when analyzing the true fill rate, the Engineering Estimate performs significantly inferior to the other methods. Thus, fill rate performance evaluation is important.

5.4.3 Fill Rate versus Forecast Accuracy

Most of the comparative studies use forecast accuracy to compare forecasting methods. Hence, this research explores if the fill rate performance and forecast accuracy results agree.

When looking at the forecast accuracy for LRU parts shown in Table 5.13, the Engineering Estimate and Current Method perform superior in Year 2. However, when looking at fill rate, Bayes' (starting the year at 28% fill rate and ending the year at 70% fill rate) performs significantly better in Year 2 than the Engineering Estimate and Current Method (starting the year at 28% fill rate and ending the year at 42% fill rate). In Year 3 and Year 4, the forecast accuracy and fill rate results agree that Bayes' performs superior to the other methods. However, when comparing the Current method in Year 3 and Year 4, the forecast accuracy results conclude that the Current Method performs worst, and in terms of fill rate, the Current Method (Year 3 \cong 42%-72%/Year 4 \cong 72%-81%) performs significantly better than the Engineering Estimate (Year 3/Year 4 \cong 39%). Table 5.22 summarizes these results and indicates what forecasting method forecasts best where (1) is best, (2) is second best, and (3) is worst. This illustrates how forecast accuracy can be a poor indicator when comparing methods, which further emphasizes the importance of evaluating performance in comparative studies.

Table 5.22 Fill Rate versus Forecast Accuracy

	Year 2	Year 3	Year 4
Fill Rate	1) Bayes 2) EE/Current	1) Bayes 2) Current 3) EE	1) Bayes 2) Current 3) EE
Forecast Accuracy	1) EE/Current 2) Bayes	1) Bayes 2) EE 3) Current	1) Bayes 2) EE 3) Current

Chapter 6 – Conclusions, Contributions and Directions for Future Research

6.1 Conclusions / Contributions

Demand forecasting plays a crucial role in inventory management. Inventory stock levels are dependent on predictions of demand, and accurately predicting spare parts demand has been an ongoing challenge in the aerospace industry, especially when aircraft programs are new. Due to the intermittent and slow-moving nature of these parts, spare parts management has received renewed attention; however, there are no conclusive results on which forecasting method is “best.” The majority of comparative studies use forecast accuracy metrics that are not suitable for data series with many zeroes, and this can lead to biased results. Thus, inventory and service levels are recommended in Teunter and Duncan (2009). Additionally, a significant amount of literature has highlighted a gap between research and practice. The theoretical knowledge has advanced; yet, the techniques developed for spare parts forecasting are usually neglected in industry due to practicality constraints such as data availability, model complexity, and forecasting support systems (Bacchetti and Saccani, 2012). Thus, this thesis contributes to existing research by proposing a practical solution for forecasting spare parts demand, especially for new aircraft programs. The proposed method was created around common data constraints faced in industry and is easy to integrate into a company’s current forecasting system. Additionally, it is important to note that this research is not intended to change how inventory is optimized. Its’ focus is to improve demand prediction allowing for more accurate computation of inventory stock levels.

The proposed approach uses Bayes’ rule to forecast demand. Bayes’ rule is a systematic approach that revises prior knowledge (engineering estimates) as observed demand occurs. This

approach is recommended when data is limited, which is the case for new programs where a majority of parts experience low demand. The proposed method includes weighted prior parameters. This gives managers the ability to incorporate their uncertainty in engineering estimates and allows for these estimates to hold importance when predicting demand. Additionally, the proposed method creates a new business rule to more accurately represent demand for certain scenarios when no demands have occurred. After developing the method, a case study was performed to validate the use of Bayes' rule when forecasting demand. This case study was performed using 3 ¼ years of data from a leading aircraft manufacturer, and three forecasting methods were analyzed: Engineering Estimates (judgmental forecasting), the Current Method (traditional statistical forecasting), and Bayes' rule (Bayesian forecasting). These methods were compared by their impact on forecast accuracy, inventory costs, and performance (fill rate in this study).

This research found that the demand forecast accuracy metrics were heavily influenced by one outlier part. Hence, it is important to be cognizant of this when performing forecast accuracy analysis on a majority of low demand parts. Additionally, when evaluating the outlier part, it was found that the engineering estimate was quickly overwhelmed by observed demand, so evaluating an appropriate weighting scheme is essential. The proposed weighting scheme for prior parameters seemed appropriate for all parts but "Off the Shelf" parts where the weights should be increased. Also, in this study, the forecast accuracy results validated that the Bayesian Method works well for parts with limited data, or otherwise defined as low demand parts in this study.

The Current Method and Bayesian Method both outperform Engineering Estimates relative to fill rate performance. Over time, both the Current Method and Bayesian Method result in

similar fill rate performance. However, Bayes' significantly outperforms the Current Method in the first few years of the program. Thus, Bayes' can be considered a first responder and is especially useful in the early life of a program.

Target stock levels and performance were most effected when the estimates of demand were first revised with observed demand. This happened in Year 2 for the Bayesian approach and Year 3 for the Current Method. In these years, the inventory investment costs significantly increased. Hence, management should absolutely budget for increased inventory investment costs on new aircraft programs as estimates of demand are refined over time.

This case study shows the importance of evaluating true fill rate performance in comparative studies, which was done using a supply chain simulation in this research. If fill rate performance analysis was not performed, the Engineering Estimate would have seemed superior when looking at the lowest inventory investment costs. However, when analyzing the true fill rate, the Engineering Estimate performs significantly inferior to the other methods. Furthermore, when looking at forecast accuracy, in Year 2 the Engineering Estimate and Current Method performs best, and the Bayesian Method performs worst. In Year 3 and Year 4, the Bayesian Method performs best, and the Current Method performs worst. However, when analyzing fill rate in Year 2, Bayes' actually performs significantly better than the Engineering Estimate and Current Method. In Year 3 and Year 4, the Current Method's performance converges with Bayes' performance, and the Current Method performs significantly better than the Engineering Estimate. Thus, forecast accuracy metrics and optimized inventory investment costs can be poor indicators of determining which method is "best." This analysis emphasizes the importance of evaluating performance in comparative studies for spare parts.

6.2 Directions for Future Research

6.2.1 Hybrid Method

Most research proposes a single forecasting model for all parts. However, In Chapter 2 it was apparent that some approaches work better for certain parts than others. This was also illustrated in the forecast accuracy results where the engineering estimate was the most effective method for “Off the Shelf” parts. Only a few papers have given criteria to differentiate forecasting approaches by different part characteristics. Kalchschmidt et al. (2003) proposed a method to filter stable and irregular demand series in order to assign these filtered groups to the appropriate method. Syntetos et al. (2005) classify parts by demand and assigns these parts to different methods based on these demand-based classifications. Croston’s method and its’ variants were used to forecast intermittent, erratic, and lumpy demand parts, and Simple Exponential Smoothing was used for “smooth” demand parts. Furthermore, Bayes’ has proven to work well when data is limited and in the early life of a program. However, another approach might prove more effective after an aircraft has reached steady state for a worthy amount of time and a significant amount of data has been accumulated. These types of hybrid methods are extremely worth investigating for the potential to further improve performance.

6.2.2 Analysis on Different Scenario

Analysis on a scenario that takes longer to reach steady state than the one analyzed in this research could further illustrate the potential of Bayes’ rule. The current aircraft analyzed took one years to reach steady state, or in other words, the total operating hours each period after Year 1 is stationary. Bayes’ rule can handle an increase in operating hours. However, Simple Exponential Smoothing (SES) is known to lag when a trend exists, which is the case as operating hours increase. The Current Method uses a blend of Simple Exponential Smoothing

(SES) and causal factors, and this blend could lead to reduced performance as operating hours increase. Additionally, the case study concludes that Bayes' and the Current Method have a similar fill rate by the third year (where steady state has already been reached). However, if steady state was not reached in Year 2, Bayes' has the potential to perform superior to the Current Method because SES is not recommended when a trend exists. Thus, further analysis on an aircraft that takes long to reach steady state could further prove the benefits of Bayes' rule relative to the Current Method.

6.2.3 Current Data Limitations

Chapter 3 discussed how the data was aggregated. In summary the data was given by aggregate operating hours per month and part order dates/quantities. This only allows for the calculation of a single parameter (average operating hours per demand) with an assumed demand distribution. Thus, this research assumed the time between failures is exponentially distributed, which is equivalent to Poisson demand (demands per time). Demand forecasts traditionally assume Poisson demand (variance equals mean). However, the literature finds that the Poisson distribution is a poor estimator of actual variance (Sherbrooke, 2004). Thus, the calculation of a demand distribution based on observed demand could determine whether this Poisson assumption needs revision. The desired data set to do this would show the total operating hours between failures for each failed asset. This could be obtained by time stamping the installation of asset Y on aircraft tail X or by time stamping the failure of asset Y on aircraft tail X. Then, the flight hours for aircraft tail X could be computed by two time stamps. This type of data set could allow for a more appropriate representation of demand and could potentially increase accuracy of predicted demand.

REFERENCES

- Abdi, H., & Williams, L. (2010). Jackknife. 655-661.
- Adrodegari, F., Bacchetti, A., Saccani, N., & Syntetos, A. (2014). Spare Parts Inventory Management. A Literature Review And Directions for Future Research.
- Armstrong, J. S., & Collopy, F. (1992). Error measures for generalizing about forecasting methods: Empirical comparisons. *International Journal of Forecasting*, 8(1), 69-80.
- Aronis, K.-P., Magou, I., Dekker, R., & Tagaras, G. (2004). Inventory control of spare parts using a Bayesian approach: A case study. *European Journal of Operational Research*, 154(3), 730-739.
- Arrow, K. J. (1958). *Studies in the mathematical theory of inventory and production*: Stanford University Press.
- Azoury, K. S., & Miller, B. L. (1984). A comparison of the optimal ordering levels of Bayesian and non-Bayesian inventory models. *Management Science*, 30(8), 993-1003.
- Bacchetti, A., & Saccani, N. (2012). Spare parts classification and demand forecasting for stock control: Investigating the gap between research and practice. *Omega*, 40(6), 722-737.
- Berkowitz, D., Gupta, J. N., Simpson, J. T., McWilliams, J., Delayne, L., Brown, B., . . . Sparks, T. (2003). Performance Based Logistics. *Center for the Management of Science and Technology, Huntsville, AL*.
- Bookbinder, J. H., & Lordahl, A. E. (1989). Estimation of inventory re-order levels using the bootstrap statistical procedure. *IIE transactions*, 21(4), 302-312.

- Boone, C. A., Craighead, C. W., & Hanna, J. B. (2008). Critical challenges of inventory management in service parts supply: A Delphi study. *Operations Management Research*, 1(1), 31-39.
- Boylan, J. E., & Syntetos, A. A. (2008). Forecasting for inventory management of service parts *Complex System Maintenance Handbook* (pp. 479-506): Springer.
- Boylan, J. E., & Syntetos, A. A. (2009). Spare parts management: a review of forecasting research and extensions. *IMA journal of management mathematics*.
- Brown, G. F., & Rogers, W. F. (1973). A Bayesian approach to demand estimation and inventory provisioning. *Naval Research Logistics Quarterly*, 20(4), 607-624.
- Brown, R. G. (1959). Statistical forecasting for inventory control.
- Callegaro, A. (2009). Forecasting Methods for Spare Parts Demand.
- Carlo, C. M. (2004). Markov Chain Monte Carlo and Gibbs Sampling.
- Cohen, M. A., Agrawal, N., & Agrawal, V. (2006). Winning in the aftermarket. *Harvard business review*, 84(5), 129.
- Congdon, P. (2003). *Applied bayesian modelling*: John Wiley & Sons.
- Croston, J. D. (1972). Forecasting and stock control for intermittent demands. *Operational Research Quarterly*, 289-303.
- De Alba, E., & Mendoza, M. (2007). Bayesian forecasting methods for short time series. *Foresight: The International Journal of Applied Forecasting*, International Institute of Forecasters(8), 41-44.

- Diamond, B., Krahl, D., Nastasi, A., & Tag, P. (2010). ExtendSim advanced technology: integrated simulation database. *Proceedings of the Winter Simulation Conference*, 32-39.
- Dolgui, A., & Pashkevich, M. (2008). Demand forecasting for multiple slow-moving items with short requests history and unequal demand variance. *International journal of production economics*, 112(2), 885-894.
- Duncan, G. T., Gorr, W. L., & Szczypula, J. (2001). Forecasting analogous time series. *Principles of forecasting* (pp. 195-213): Springer.
- Eaves, A., & Kingsman, B. (2004). Forecasting for the ordering and stock-holding of spare parts. *Journal of the Operational Research Society*, 55(4), 431-437.
- Efron, B. (1979). Bootstrap methods: another look at the jackknife. *The annals of Statistics*, 1-26.
- Gelfand, A. E., & Smith, A. F. (1990). Sampling-based approaches to calculating marginal densities. *Journal of the American statistical association*, 85(410), 398-409.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis*: CRC press.
- Ghobbar, A. A., & Friend, C. H. (2003). Evaluation of forecasting methods for intermittent parts demand in the field of aviation: a predictive model. *Computers & Operations Research*, 30(14), 2097-2114.
- Goodwin, P. (2002). Integrating management judgment and statistical methods to improve short-term forecasts. *Omega*, 30(2), 127-135.

- Grange, F. (1998). Challenges in modeling demand for inventory optimization of slow-moving items *Proceedings of the 30th conference on Winter simulation* (pp. 1211-1218): IEEE Computer Society Press.
- Graves, S. C., Kan, A. R., & Zipkin, P. H. (1993). *Logistics of production and inventory* (Vol. 4): Elsevier.
- Gross, O. (1956). A class of discrete-type minimization problems: DTIC Document.
- Gutierrez, R. S., Solis, A. O., & Mukhopadhyay, S. (2008). Lumpy demand forecasting using neural networks. *International journal of production economics*, *111*(2), 409-420.
- Hayes, R. H. (1969). Statistical estimation problems in inventory control. *Management Science*, *15*(11), 686-701.
- Hill, R. M. (1999). Bayesian decision-making in inventory modelling. *IMA journal of management mathematics*, *10*(2), 147-163.
- Hill, T., O'Connor, M., & Remus, W. (1996). Neural network models for time series forecasts. *Management Science*, *42*(7), 1082-1092.
- Holt, C. (1957). Forecasting trends and seasonal by exponentially weighted moving averages. *ONR Memorandum*, *52*.
- Iglehart, D. L. (1964). The dynamic inventory problem with unknown demand distribution. *Management Science*, *10*(3), 429-440.
- Johnston, F., & Boylan, J. E. (1996). Forecasting for items with intermittent demand. *Journal of the Operational Research Society*, 113-121.

- Kalchschmidt, M., Zotteri, G., & Verganti, R. (2003). Inventory management in a multi-echelon spare parts supply chain. *International journal of production economics*, 81, 397-413.
- Kamath, K. R., & Pakkala, T. (2002). A Bayesian approach to a dynamic inventory model under an unknown demand distribution. *Computers & Operations Research*, 29(4), 403-422.
- Kaplan, A. J. (1976). Inventory problem with unknown mean demand and learning. *Naval Research Logistics Quarterly*, 23(4), 687-695.
- Kaplan, A. J. (1988). Bayesian approach to inventory control of new parts. *IIE transactions*, 20(2), 151-156.
- Klassen, R. D., & Flores, B. E. (2001). Forecasting practices of Canadian firms: Survey results and comparisons. *International journal of production economics*, 70(2), 163-174.
- Kolassa, S., & Schütz, W. (2007). Advantages of the MAD/MEAN ratio over the MAPE. *Foresight: The International Journal of Applied Forecasting*(6), 40-43.
- Levén, E., & Segerstedt, A. (2004). Inventory control with a modified Croston procedure and Erlang distribution. *International journal of production economics*, 90(3), 361-367.
- Li, S., & Kuo, X. (2008). The inventory management system for automobile spare parts in a central warehouse. *Expert Systems with Applications*, 34(2), 1144-1153.
- Li, S., Ragu-Nathan, B., Ragu-Nathan, T. S., & Subba Rao, S. (2006). The impact of supply chain management practices on competitive advantage and organizational performance. *Omega*, 34(2), 107-124.

- Li, X. (2014). Operations Management of Logistics and Supply Chain: Issues and Directions. *Discrete Dynamics in Nature and Society*, 2014, 7. doi: 10.1155/2014/701938
- Louit, D., Pascual, R., Banjevic, D., & Jardine, A. K. S. (2011). Optimization models for critical spare parts inventories—a reliability approach. *J Oper Res Soc*, 62(6), 992-1004.
- McCarthy, T. M., Davis, D. F., Golicic, S. L., & Mentzer, J. T. (2006). The evolution of sales forecasting management: a 20-year longitudinal study of forecasting practices. *Journal of Forecasting*, 25(5), 303-324.
- McGrayne, S. B. (2011). *The theory that would not die: how Bayes' rule cracked the enigma code, hunted down Russian submarines, & emerged triumphant from two centuries of controversy*: Yale University Press.
- Muñoz, D. F., & Muñoz, D. F. (2011). Bayesian forecasting of spare parts using simulation *Service Parts Management* (pp. 105-123): Springer.
- O'Hagan, A., & Luce, B. (2003). A primer on Bayesian statistics in health economics and outcomes research. *Bethesda, MD: MEDTAP International*.
- Petrović, R., Šenborn, A., & Vujošević, M. (1989). A new adaptive algorithm for determination of stocks in spare parts inventory systems. *Engineering Costs and Production Economics*, 15, 405-410.
- Porras Musalem, E. (2005). Inventory theory in practice: Joint replenishments and spare parts control (Chapter 6). *PhD Thesis, Erasmus University Rotterdam (Tinbergen Institute), The Netherlands*.

- Rahman, M. A., & Sarker, B. R. (2012). A Bayesian approach to forecast intermittent demand for seasonal products. *International Journal of Industrial and Systems Engineering*, 11(1), 137-153.
- Recchia, C. (2012). Bayesian Methods in Reliability Engineering.
- Sani, B., & Kingsman, B. (1997). Selecting the best periodic inventory control and demand forecasting methods for low demand items. *Journal of the Operational Research Society*, 700-713.
- Saylor, S. E., & Dailey, J. K. (2010). *Advanced logistics analysis capabilities environment*. Paper presented at the Simulation Conference (WSC), Proceedings of the 2010 Winter.
- Scarf, H. (1959). Bayes solutions of the statistical inventory problem. *The Annals of Mathematical Statistics*, 490-508.
- Scarf, H. E. (1960). Some remarks on Bayes solutions to the inventory problem. *Naval Research Logistics Quarterly*, 7(4), 591-596.
- Sherbrooke, C. C. (2004). *Optimal inventory modeling of systems: multi-echelon techniques* (Vol. 72): Springer.
- Soliman, A. A., Abd Allah, A. H., & Sultan, K. (2006). Comparison of estimates using record statistics from Weibull model: Bayesian and non-Bayesian approaches. *Computational Statistics & Data Analysis*, 51(3), 2065-2077.
- Syntetos, A., Boylan, J., & Croston, J. (2005). On the categorization of demand patterns. *Journal of the Operational Research Society*, 56(5), 495-503.

- Syntetos, A. A., & Boylan, J. E. (2001). On the bias of intermittent demand estimates. *International journal of production economics*, 71(1), 457-466.
- Syntetos, A. A., & Boylan, J. E. (2005). The accuracy of intermittent demand estimates. *International Journal of Forecasting*, 21(2), 303-314.
- Syntetos, A. A., Boylan, J. E., & Disney, S. M. (2009). Forecasting for inventory planning: a 50-year review. *Journal of the Operational Research Society*, S149-S160.
- Syntetos, A. A., Keyes, M., & Babai, M. (2009). Demand categorisation in a European spare parts logistics network. *International Journal of Operations & Production Management*, 29(3), 292-316.
- Syntetos, A. A., Nikolopoulos, K., Boylan, J. E., Fildes, R., & Goodwin, P. (2009). The effects of integrating management judgement into intermittent demand forecasts. *International journal of production economics*, 118(1), 72-81.
- Teunter, R. H., & Duncan, L. (2009). Forecasting intermittent demand: a comparative study. *Journal of the Operational Research Society*, 60(3), 321-329.
- Ton de, K., Janssen, F., Van Doremalen, J., Van Wachem, E., Clerkx, M., & Peeters, W. (2005). Philips Electronics Synchronizes Its Supply Chain to End the Bullwhip Effect. *Interfaces*, 35(1), 37-48.
- Tukey, J. W. (1958). *Bias and confidence in not-quite large samples*. Paper presented at the Annals of Mathematical Statistics.
- Turner, D. (1990). The role of judgement in macroeconomic forecasting. *Journal of Forecasting*, 9(4), 315-345.

- Wagner, S., & Lindemann, E. (2008). A case study-based analysis of spare parts management in the engineering industry. *Production planning & control*, 19(4), 397-407.
- Wang, Z. F., & Mersereau, A. J. (2013). Bayesian Inventory Management with Potential Change-Points in Demand. *Working Paper, University of North Carolina*.
- Willemain, T. R., Smart, C. N., & Schwarz, H. F. (2004). A new approach to forecasting intermittent demand for service parts inventories. *International Journal of Forecasting*, 20(3), 375-387.
- Willemain, T. R., Smart, C. N., Shockor, J. H., & DeSautels, P. A. (1994). Forecasting intermittent demand in manufacturing: a comparative evaluation of Croston's method. *International Journal of Forecasting*, 10(4), 529-538.
- Winters, P. R. (1960). Forecasting sales by exponentially weighted moving averages. *Management Science*, 6(3), 324-342.
- Yelland, P. M. (2010). Bayesian forecasting of parts demand. *International Journal of Forecasting*, 26(2), 374-396.
- Zacks, S. (1969). Bayes sequential design of stock levels. *Naval Research Logistics Quarterly*, 16(2), 143-155.