

RELIABLE PREDICTION INTERVALS AND BAYESIAN ESTIMATION  
FOR DEMAND RATES OF SLOW-MOVING INVENTORY

Matthew Douglas Lindsey, B.S., M.B.A.

Dissertation Prepared for the Degree of  
DOCTOR OF PHILOSOPHY

UNIVERSITY OF NORTH TEXAS

August 2007

APPROVED:

Robert Pavur, Major Professor  
M. Theodore Farris, Committee Member  
Shailesh Kulkarni, Committee Member  
Ching-Chung Kuo, Committee Member  
Sherry Ryan, Advisor for Department of  
Information Technology and Decision  
Science  
Mary Jones, Chair of the Department of  
Information Technology and Decision  
Science  
Kathleen B. Cooper, Dean of the College of  
Business Administration  
Sandra L. Terrell, Dean of the Robert B.  
Toulouse School of Graduate Studies

Lindsey, Matthew Douglas. Reliable Prediction Intervals and Bayesian Estimation for Demand Rates of Slow-Moving Inventory. Doctor of Philosophy (Management Science), August 2007, 180 pp., 58 figures, 17 tables, references, 122 titles.

Inventory having intermittent demand has infrequent sales that appear at random, with many periods that do not show any demand at all. Managing inventory with intermittent demand has received less attention in the literature than that of fast-moving products. This is due in part, perhaps, to the lack of observable historical sales figures for inventory with intermittent demand or because slow-moving inventory does not provide the bulk of sales, despite often being the bulk of inventory on hand.

Inventory management tools are proposed that provide estimation procedures for the future demand rates of inventory with intermittent demand. Prediction intervals, adapted from statistical procedures developed for software reliability, for the future demand rate of a group of products that have no sales or no more than one sale over a specified time frame are proposed. A Monte Carlo simulation study is conducted to assess the reliability of these prediction intervals across various sizes of product groups and demand rates as well as for mixtures of demand rates and identify reliable parameter ranges. Sales data from a Fortune 500 company were used to assess the performance of the proposed prediction intervals.

Inventory managers periodically update their predictions of future demand rates for products. Two models – a Bayes model, using a prior probability distribution for the demand rate and a Poisson model, using a Poisson distribution for demand – were used to obtain optimal inventory levels over several periods assuming a known cost for surplus and shortage. This procedure has been proposed in the literature. However, its

performance has not been examined under various demand rates such as intermittent demand.

A Monte Carlo simulation study was used to examine the performance of the Bayes and Poisson model under moderate and intermittent demand. When the demand rates of the products are homogeneous, the inventory costs related to the Bayes model is lower than that of the Poisson model. The Poisson model is preferred under conditions of high variability among product demand rates. An improvement that optimized inventory costs for some demand rates was made to the Bayes model using a mixture of priors.

Copyright 2007

by

Matthew Douglas Lindsey

## ACKNOWLEDGMENTS

I am indebted to many people who tirelessly helped me through the Ph.D. program at the University of North Texas by offering knowledge, advice, camaraderie, and support. Although this work bears my name, it would not have been possible without the help of many individuals in the UNT College of Business, including the faculty, staff and fellow PhD students, past and present, as well as my colleagues at The University of Texas at Tyler. I am grateful to all who played a role. I especially would like to thank my dissertation chair and committee who have all worked diligently with me on this endeavor. I am lucky to have had the opportunity to work with this team of researchers. My dissertation chair has truly become a good friend and mentor and helped with learning statistical concepts as well as working in the academic environment.

I would also like to thank my friends, family, and former coworkers for being supportive of my decision to return to school and pursue my degree. I am lucky to have had the opportunity to take advantage of the education offered at the University of North Texas by the citizens of the great state of Texas.

## TABLE OF CONTENTS

ACKNOWLEDGMENTS .....	iii
LIST OF TABLES.....	vi
LIST OF FIGURES .....	viii
CHAPTER	
1. INTRODUCTION .....	1
Computerized Inventory Management Systems .....	3
Controlling Inventory .....	4
Definition and Examples of Slow-Moving Inventory .....	5
Methodologies for Estimating Demands for Slow-Moving Products .....	8
Significance of the Study.....	11
2. LITERATURE REVIEW .....	14
Inventory Models.....	14
Periodic and Continuous Replenishment Models.....	15
Continuous Review: (s, S) Model and (s, Q) Model .....	16
Assumptions of (s, S) Model .....	16
Assumptions of (s, Q) Model.....	17
Assumptions of (S - 1, S) Model.....	17
Inventory Models for Spare Parts.....	18
Inventory Model Cost Structures .....	21
Demand for Slow-Moving Inventory .....	22
Forecasting Methods.....	23
Lumpy Demand.....	25
Most Recent Research Related to Slow-Moving Inventory .....	28
Demand Distributions.....	28
Poisson Distribution of Demand for Slow-Moving Inventory.....	29
Compound Poisson Distribution of Demand for Slow-Moving Inventory .....	30
Gamma Distribution of Demand for Slow-Moving Inventory.....	32
Normal Distribution of Demand for Slow-Moving Inventory .....	33
Hermite Distribution of Demand for Slow-Moving Inventory.....	34
Other Distributions of Demand for Slow-Moving Inventory.....	35
Bayesian Approach to Inventory Modeling.....	36
Software Reliability.....	37
Decision Making Involving Slow-Moving Inventory.....	41

3. METHODOLOGY .....	45
Proposed Methodologies .....	45
Estimating Demand Rate for a Single Product with Nonzero Demand.....	47
Estimating Demand Rate for a Family of Products with Zero Demand .....	49
Bayesian Methodology for Inventory Modeling .....	59
Stopping Rules.....	62
Research Questions.....	64
4. ANALYSIS AND RESULTS .....	68
Research Question 1 .....	69
Effect of Product Group Size on Prediction Intervals .....	70
Effect of Demand Rates on Two-Sided Prediction Intervals.....	76
Two-Sided Prediction Intervals Using a Mixture of Demand Rates .....	78
Research Question 2 .....	85
Research Question 3 .....	90
Research Question 4 .....	100
Determining Optimal Inventory Levels Based on Costs and Distribution .....	102
Research Question 5 .....	113
Research Question 6 .....	127
Research Question 7 .....	134
5. DISCUSSION AND IMPLICATIONS.....	143
Discussion of Research Question 1 .....	143
Discussion of Research Question 2 .....	148
Discussion of Research Question 3 .....	150
Discussion of Research Questions 4 and 5 .....	153
Discussion of Research Question 6 .....	155
Discussion of Research Question 7 .....	157
Managerial Implications .....	159
Limitations .....	166
Unique Contribution of Research and Future Research Ideas.....	168
REFERENCES .....	171

## LIST OF TABLES

Table 1 Characteristics of Different Types of Demand.....	2
Table 2 Product Examples Based on Replenishment Period and Demand Size.....	8
Table 3 Forecasting Methodologies for Different Demand Sizes and Replenishment Periods.....	10
Table 4 Research on Spare Parts Inventory .....	20
Table 5 Demand Forecasting Methodologies .....	25
Table 6 Demand Distributions Cited in Literature .....	36
Table 7 Empirical Type I Error for TSPI in Figure 11 .....	73
Table 8 Empirical Type I Error for Zero and One Sales OSPI and TSPI in <i>Figure 23</i> .....	99
Table 9 Comparison of Bayes Model with Poisson Model Ranked by Coefficient of Variation (CV) .....	125
Table 10 Comparison of Bayes/Ross Model with Poisson Model Ranked by ation Comparison of Bayes/Ross Model with Poisson Model Ranked by Coefficient of Variation (CV) .....	129
Table 11 Summary of Reliable Prediction Intervals across Retail Stores .....	137
Table 12 Numbers and Percentages of Reliable Prediction Intervals across 30 Stores.....	138
Table 13 Number of Future Demand Rates within each Interval across Product Families umber of future demand rates within each interval across product families .....	139

Table 14 Three-Way ANOVA of the Factors Affecting Reliability of Prediction Intervals .....	140
Table 15 Results from Tukey and Bonferroni Tests at 5% Significance Level.....	141
Table 16 Mixtures of MTBDs and Reliability for 95% Prediction Intervals.....	147
Table 17 Reliability of Prediction Intervals for Products of Sizes of 200 and 600.....	151

## LIST OF FIGURES

<i>Figure 1.</i> Distribution with no sales over a period of 40 hours. ....	57
<i>Figure 2.</i> Distribution with no sales over a period of 70 hours. ....	57
<i>Figure 3.</i> Distribution with no sales over a period of 130 hours. ....	58
<i>Figure 4.</i> Distribution with no sales over a period of 150 hours. ....	58
<i>Figure 5.</i> Empirical Type I error versus product group size for TSPI.....	72
<i>Figure 6.</i> Empirical Type I error versus MTBD for TSPI for 200 products. ....	77
<i>Figure 7.</i> Empirical Type I error for TSPI for mixtures of MTBDs with MTBD of 10 for second subgroup. ....	80
<i>Figure 8.</i> Empirical Type I error for TSPI for mixture of MTBDs with MTBD of 50 for second subgroup. ....	81
<i>Figure 9.</i> Empirical Type I error for TSPI for mixture of MTBDs with MTBD of 100 for second subgroup. ....	82
<i>Figure 10.</i> Empirical Type I error for TSPI for mixture of MTBDs with MTBD of 400 for second subgroup. ....	83
<i>Figure 11.</i> Empirical Type I error versus product group size for OSPI.....	88
<i>Figure 12.</i> Empirical Type I error versus MTBD for OSPI for 200 products. ....	89
<i>Figure 13.</i> Empirical Type I error for zero and one sales and zero sales OSPI. ....	92
<i>Figure 14.</i> Empirical Type I error for zero and one sales and zero sales TSPI.....	93
<i>Figure 15.</i> Empirical Type I error for zero and one sales OSPI across product group size. ....	95

<i>Figure 16.</i> Empirical Type I error for zero and one sales TSPI across product group size. ....	96
<i>Figure 17.</i> Empirical Type I error for zero and one sales OSPI and TSPI.....	98
<i>Figure 18.</i> Cost per time period when demand mean is 3, standard deviation is 5.5 and surplus to shortage cost ratio is 1 to 5. ....	105
<i>Figure 19.</i> Cost per time period when demand mean is 3, standard deviation is 5.5 and surplus to shortage cost ratio is 1 to 1. ....	106
<i>Figure 20.</i> Cost per time period when demand mean is 3, standard deviation is 5.5 and surplus to shortage cost ratio is 5 to 1. ....	106
<i>Figure 21.</i> Cost per time period when demand mean is 3, standard deviation is 1.7 and surplus to shortage cost ratio is 1 to 5. ....	107
<i>Figure 22.</i> Cost per time period when demand mean is 3, standard deviation is 1.7 and surplus to shortage cost ratio is 1 to 1. ....	108
<i>Figure 23.</i> Cost per time period when demand mean is 3, standard deviation is 1.7 and surplus to shortage cost ratio is 5 to 1. ....	108
<i>Figure 24.</i> Cost per time period when demand mean is 3, standard deviation is 0.55 and surplus to shortage cost ratio is 1 to 5. ....	109
<i>Figure 25.</i> Cost per time period when demand mean is 3, standard deviation is 0.55 and surplus to shortage cost ratio is 1 to 1. ....	110
<i>Figure 26.</i> Cost per time period when demand mean is 3, standard deviation is 0.55 and surplus to shortage cost ratio is 5 to 1. ....	110
<i>Figure 27.</i> Cost per time period when demand mean is 3, standard deviation is 0.17 and surplus to shortage cost ratio is 1 to 5. ....	111

<i>Figure 28.</i> Cost per time period when demand mean is 3, standard deviation is 0.17 and surplus to shortage cost ratio is 1 to 1. ....	112
<i>Figure 29.</i> Cost per time period when demand mean is 3, standard deviation is 0.17 and surplus to shortage cost ratio is 5 to 1. ....	112
<i>Figure 30.</i> Cost per time period when demand mean is 0.1, standard deviation is 0.0316 and surplus to shortage cost ratio is 1 to 5.....	115
<i>Figure 31.</i> Cost per time period when demand mean is 0.1, standard deviation is 0.3162 and surplus to shortage cost ratio is 1 to 5.....	115
<i>Figure 32.</i> Cost per time period when demand mean is 0.1, standard deviation is 1.0 and surplus to shortage cost ratio is 1 to 5. ....	116
<i>Figure 33.</i> Cost per time period when demand mean is 0.05, standard deviation is 0.0071 and surplus to shortage cost ratio is 1 to 5.....	116
<i>Figure 34.</i> Cost per time period when demand mean is 0.05, standard deviation is 0.0224 and surplus to shortage cost ratio is 1 to 5.....	117
<i>Figure 35.</i> Cost per time period when demand mean is 0.05, standard deviation is 0.0707 and surplus to shortage cost ratio is 1 to 5.....	117
<i>Figure 36.</i> Cost per time period when demand mean is 0.05, standard deviation is 0.2236 and surplus to shortage cost ratio is 1 to 5.....	118
<i>Figure 37.</i> Cost per time period when demand mean is 0.01, standard deviation is 0.0032 and surplus to shortage cost ratio is 1 to 5.....	118
<i>Figure 38.</i> Cost per time period when demand mean is 0.01, standard deviation is 0.01 and surplus to shortage cost ratio is 1 to 5. ....	119

*Figure 39.* Cost per time period when demand mean is 0.01, standard deviation is 0.0316 and surplus to shortage cost ratio is 1 to 5..... 119

*Figure 40.* Cost per time period when demand mean is 0.01, standard deviation is 0.1 and surplus to shortage cost ratio is 1 to 5. .... 120

*Figure 41.* Cost per time period when demand mean is 0.005, standard deviation is 0.0022 and surplus to shortage cost ratio is 1 to 5..... 120

*Figure 42.* Cost per time period when demand mean is 0.005, standard deviation is 0.0071 and surplus to shortage cost ratio is 1 to 5..... 121

*Figure 43.* Cost per time period when demand mean is 0.005, standard deviation is 0.0224 and surplus to shortage cost ratio is 1 to 5..... 121

*Figure 44.* Cost per time period when demand mean is 1/300, standard deviation is 0.0005 and surplus to shortage cost ratio is 1 to 5..... 122

*Figure 45.* Cost per time period when demand mean is 1/300, standard deviation is 0.0058 and surplus to shortage cost ratio is 1 to 5..... 122

*Figure 46.* Cost per time period when demand mean is 0.002, standard deviation is 0.0014 and surplus to shortage cost ratio is 1 to 5..... 123

*Figure 47.* Cost per time period when demand mean is 0.002, standard deviation is 0.0141 and surplus to shortage cost ratio is 1 to 5..... 123

*Figure 48.* Cost per time period when demand mean is 0.001, standard deviation is 0.001 and surplus to shortage cost ratio is 1 to 5. .... 124

*Figure 49.* Bayes/Ross Model versus Poisson Model, demand mean is 0.1 and standard deviation is 3.162 with surplus to shortage cost ratio of 1 to 5..... 131

<i>Figure 50.</i> Bayes/Ross Model versus Poisson Model, demand mean is 0.01 and standard deviation is 0.3162 with surplus to shortage cost ratio of 1 to 5.....	131
<i>Figure 51.</i> Bayes/Ross Model versus Poisson Model, demand mean is 0.002 and standard deviation is 0.0447 with surplus to shortage cost ratio of 1 to 5.....	132
<i>Figure 52.</i> Bayes/Ross Model versus Poisson Model, demand mean is 1/300 and standard deviation is 0.0577 with surplus to shortage cost ratio of 1 to 5.....	132
<i>Figure 53.</i> Bayes/Ross Model versus Poisson Model, demand mean is 0.005 and standard deviation is 0.0707 with surplus to shortage cost ratio of 1 to 5.....	133
<i>Figure 54.</i> Bayes/Ross Model versus Poisson Model, demand mean is 0.05 and standard deviation is 0.2236 with surplus to shortage cost ratio of 1 to 5.....	133
<i>Figure 55.</i> Bayes/Ross Model versus Poisson Model, demand mean is 0.1 and standard deviation is 0.3162 with surplus to shortage cost ratio of 1 to 5.....	134
<i>Figure 56.</i> Mean of factor levels versus confidence level.....	142
<i>Figure 57.</i> Usefulness of proposed prediction intervals and models. ....	164
<i>Figure 58.</i> General conditions and assumptions for proposed prediction intervals and models.....	165

# CHAPTER 1

## INTRODUCTION

Inventory, or a stock of goods, is a necessary part of business and comes in many forms, from raw materials to work in process to finished goods. Inventory serves many functions for organizations. At the retail level, its main function is to allow the firm to meet expected customer demand and prevent shortages. However, holding inventory presents real costs to companies and thereby, managers are under pressure to reduce inventories and associated expenses (Masters, 1993). Techniques to minimize inventory while still providing high customer satisfaction have been widely studied with relatively little focus on those for products with low demand (Hollier, Mak, & Lai, 2002).

Depending on the variability of its size and the percentage of the overall demand it accounts for, demand for products may be classified into one of the four categories shown in Table 1. The demand for products in the High Demand Size Variability category of Table 1, such as trendy popular products and expensive merchandise, may be difficult to forecast because of their sporadic or intermittent nature. This research is motivated by a renewed interest in forecasting demand for retail products that are often characterized by infrequent transactions (Willemain, Smart, Shockor, & DeSautels, 1994; Johnston & Boylan, 1996; Syntetos & Boylan, 2001). In particular, we will propose methodologies for predicting demand rates for intermittent or slow-moving merchandise that are adapted from statistical procedures developed for software reliability or for Bayesian estimation.

Table 1

Characteristics of Different Types of Demand

Demand Size Variability	Percentage Accounting for 80% of Overall Demand	Percentage Accounting for 20% of Overall Demand
High	Erratic buying of low cost or trendy popular products	Intermittent sales of expensive purchases or accessories
Low	Standard products or necessities with predictable demand	Custom, seasonal, or periodic orders

Even today, consumers have access to an unprecedented variety of goods and services due to fierce global competition, shrinking product development cycles and increasing manufacturing flexibility (Fisher, Hammond, Obermeyer, & Raman, 1994). In many industries, this has resulted in a plethora of products to manage in the supply chain. Consequently, it is not unusual for retailers to have thousands of stock keeping units (SKUs). For instance, Saks Fifth Avenue of Saks, Inc. is reported to carry 400,000 to 500,000 SKUs representing from 8,000 to 9,000 styles, with 60% of them only providing 20% of the sales and 15% of them providing 50% of the sales (Gentry, 2003). This example clearly demonstrates the Pareto principle.

Also called the 80:20 Rule, The Law of the Unequal Distribution of Results or A Few Account for Most, the Pareto principle states that in many businesses a few products generate the majority of sales while the majority of products produce few sales. Predicting the demand rates for slow-moving merchandise, which accounts for 20% of the overall demand, is typically more difficult than that for products with strong demand rates. In practice, the percentage of the demand for slow-moving products will vary from the 20% value displayed in Figure 1. The

Pareto principle is evident in many case studies and suggests that a large proportion of a stock keeper's products will have low and infrequent demand (Johnston, Boylan, & Shale, 2003). Fisher et al. (1994) note that as retail product choices increase, so does the difficulty of forecasting demand and of planning production orders, resulting in an increase in forecasting errors and relevant related costs. Miragliotta and Staudacher (2004) argue that continuing product innovation has deteriorated classic, predictable demand patterns as lead times are reduced in response to the higher degree of product customization.

### Computerized Inventory Management Systems

To cope with the challenges of efficiently controlling inventory, companies have invested heavily in computer software packages such as enterprise resource planning (ERP) systems. The ERP software generally contains a material procurement module that handles the inventory management function by calculating the safety stock and reorder point based on a product's demand history (Razi & Tarn, 2003). ERP builds on manufacturing resource planning (MRP II) and material requirements planning (MRP) that translate product demand into planning schedules. Add-ins are available to supplement its existing capabilities. These integrated modules enhance the software's ability to forecast demand for fast-moving products, but often they perform poorly for products with low demand (Razi & Tarn, 2003).

## Controlling Inventory

Since the development of the just-in-time (JIT) approach at Toyota Motor Company of Japan by Taiichi Ohno, firms adopting the tenets of JIT have shown improved performance by eliminating waste and excess inventory (Celley, Clegg, Smith, & Vonderembse, 1986). Although JIT efforts were predominantly concerned with manufacturing, the awareness of inventory cost has propagated throughout the supply chain. Recently, in an attempt to become leaner, managers are attempting to decrease inventories at all levels without lowering service levels (Chopra, Reinhardt, & Dada, 2004).

Numerous research articles have investigated the applications of mathematical models to optimize inventory under a variety of conditions, such as random lead times, specific service levels, and varying product demands (Bagchi, Hayya, & Ord, 1983; Lee & Nahmias, 1993; Tyworth, Guo, & Ganeshan, 1996; Hollier et al., 2002; Chopra et al., 2004). Considering the plethora of research on inventory management, relatively little attention has been paid to products with sporadic or low demand (Hollier, Mak, & Lam, 1995). Further, much of the limited interest in low-demand products has been concerned with stocking decisions on military and industrial spare parts (Smith & Vemuganti, 1969; Haber & Sitgreaves, 1970; Bartakke, 1981; Popovic, 1987; Weingart, 1991; Razi & Tarn, 2003; Eaves & Kingsman, 2004; Dolgui & Pashkevich, 2007).

Several research articles propose analytical models for managing slow-moving inventory by assuming various demand distributions (Silver, 1965, 1991; Croston, 1972; Bagchi et al., 1983; Schultz, 1989; Haddock, Iyer, & Nagar, 1994;

Chang, Chung, & Yang, 2001; Grahovac & Chakravarty, 2001). However, none of them has been shown to be superior to the others when benchmarked against simulated or actual data (Gelders & Groenweghe, 1985; Price & Haynsworth, 1986; Willemain et al., 1994; Johnston & Boylan, 1996; Levén & Segerstedt, 2004; Ghobbar, 2004). While the behavior of products with low demand has been examined extensively, this topic is still fertile ground for research (Snyder, 2002).

Management's desire to reign in inventory costs is not limited to work in process or finished goods. Many studies describe similar challenges facing large manufacturers who must efficiently control an inventory of tens of thousands of maintenance and replacement parts (Ward, 1978; Gelders & Van Looy, 1978; Dunsmuir & Snyder, 1989). As in the retail industry, a few parts often represent the bulk of the investment and the majority of the demand in the manufacturing sector of the economy.

If a product has not sold over a specified duration of time, its demand would be projected to be zero based on any of the popular forecasting models. Yet, this product may still sell enough in the future to be worth carrying particularly if the inventory cost is well managed. This study examines the demand for these types of products and develops several methodologies to address related issues.

#### Definition and Examples of Slow-Moving Inventory

Levén and Segerstedt (2004) state that intermittent demand seems to appear at random and there are many periods, i.e., production days, weeks or even months that do not show any demand at all. Intermittent or sporadic are

other terms used in the literature to denote slow moving. A product can be defined as intermittent if it has zero demand in several periods (Segerstedt, 1994). Razi and Tarn (2003) define items with demand of more than zero and less than or equal to 50 units over a two-year period as being slow moving. Johnston et al. (2003) mention a demand rate of 12 transactions per year as being slow moving. Slow moving is not related to the dollar value associated with the demand of a product.

In the literature, the interpretation of the term slow-moving product varies according to different applications. For example, the term lumpy demand is applied to the demand of slow-moving products demonstrating especially strong sales variation over time with some extremely slow periods (Miragliotta & Staudacher, 2004). A product with intermittent demand has zero sales for many time periods. The term intermittent demand is not consistently defined in the literature and could refer to large demand with long periods between demands. In this research, slow-moving demand refers to demand characterized by periods of zero demand and low demand. Thus, slow-moving products will be defined as products having intermittent demand, resulting in long periods between demands, with small demands when they occur. For products assumed to be slow moving for this study, the average number of time periods between demands is 20% or more of the specified time frame. This definition is motivated by the performance of proposed prediction intervals in this dissertation for estimating demand rates of slow-moving products.

Boylan, Syntetos, and Karakostas (2007) provide the definitions for terms used in the study of slow-moving inventory. Products with infrequent demand occurrences are classified as intermittent. Slow-moving items have low average demand. If the demand size is highly variable, it can be classified as erratic demand. When demand is intermittent with high variations, the term lumpy is used. Finally clumped demand is used when the demand is varied between intermittent and near constant demand.

Many retail items sell slowly, including furniture, suits, high-end televisions, maintenance and replacement parts, and others that are not part of the main theme of a store such as sports clothing in a shoe store. Estimating future demand for these products may not be easy after a few weeks of sporadic sales. Harrington (2003) points out the difficulty at Bulgari in forecasting demand for its slow-moving lines of high-end fashion products such as silks and watches. Furthermore, Snyder (2002) discusses a similar challenge for automobile parts.

Masters (1993) cites several examples of products that have low demand, ranging from clothing items to automobile repair parts to compact disc titles. In particular, he describes a certain blue oxford cloth, button-down collar shirt with specific neck and sleeve sizes that might have no demand or a very low demand in a given time period despite normal demand for similar products. Haber and Sitgreaves (1970) present their work on stocking repair parts on naval vessels and illustrate a situation in which many slow-moving items have no demand. A pooling procedure was proposed to assist in forecasting the demand for items with zero usage.

Table 2 provides examples of products that are categorized according to replenishment period and demand size. In general, retail products tend to advance through a life cycle of introduction, growth, maturity, and finally decline. Products in the introduction or decline stage may be considered slow moving.

Table 2

Product Examples Based on Replenishment Period and Demand Size

Replenishment Period	Demand	
	Low	High
Short	Electrical equipment	Food, produce
Long	Aircraft equipment, designer watches, spare parts	Consumer goods

#### Methodologies for Estimating Demands for Slow-Moving Products

Most of the existing approaches to estimating demands for slow-moving products find their roots in research related to predicting usage rates for military spare parts, especially those onboard ships (Haber & Sitgreaves, 1970). Haber and Sitgreaves (1970) survey several forecasting methods for goods with sporadic demand patterns, including a widely practiced one that relies on expert opinion and is conservative, resulting in more inventory than is actually necessary.

Exponential smoothing is a quantitative methodology commonly used by manufacturers but its application requires a demand history. The technique proposed by Haber and Sitgreaves (1970) is based on usage information about a

class of parts. It is useful for predicting demand when some parts have zero sales over a specified time period. A limitation is that parts need to be classified into appropriate categories, which might not be indisputably superior to lumping products together. Another disadvantage of this method is the tendency to consistently over or under estimate future demand for individual products. While overall inventory costs seem to be slightly lower by following this approach, the risk of supply shortages is not decreased.

Inventory control policy is most often developed to strike a sound balance between the shortage costs and the costs to maintain a required service level. Silver (1965) proposes that inventory levels be first based on the desired level of service and then updated according on historical usage. Smith and Vemuganti (1969) suggest that Silver's (1965) method be modified to focus on when replacement parts are required instead of service level. Haddock et al. (1994) present a heuristic that orders up to the desired level as indicated by changing demand patterns.

Johnston and Boylan (1996) identify exponentially weighted moving average (EWMA) coupled with the mean absolute deviation (MAD) of the forecast errors as a viable method to estimate demand in a variety of inventory control situations. However, they suggest that the accuracy of the EWMA estimates for slow-moving products are poor and can be improved by using a separate model based on inter-order intervals. In further investigations, Johnston et al. (2003) find that products with high demand often have a higher number of its items sold in a single purchase, which may affect the performance of some

forecasting models. Table 3 summarizes forecasting methodologies frequently used in practice depending on the demand size and replenishment period.

As Table 3 indicates, Croston’s method (Croston, 1972, 1974) has been recommended for predicting the demand of slow-moving products that have long replenishment periods and thus few observable sales periods. Snyder (2002) compares Croston’s method with simple exponential smoothing and introduces two variations of the former: the log-space adaptation and the adaptive variance version. Both variations attempt to deal with negative demands by introducing additional parameters and can be applied to either slow-moving or fast-moving time series. Bootstrapping (Willemain, Smart, & Schwarz, 2004) has emerged as a relatively new approach for forecasting intermittent demand and it outperforms Croston’s method as well as exponential smoothing on several large data sets from industry.

Table 3

Forecasting Methodologies for Different Demand Sizes and Replenishment Periods

Replenishment Period	Demand	
	Low	High
Short	Exponential smoothing	Smoothing techniques, moving averages
Long	Croston’s method, exponential smoothing	Moving averages, ARIMA models, order overplanning

Table 3 lists order overplanning (Bartezzaghi, Verganti, & Zotteri, 1999b) and autoregressive integrated moving average (ARIMA) models (Bartezzaghi et al., 1999b) as useful for high demand products with long replenishment periods. Order overplanning is a set of techniques suitable to master production scheduling for dealing with the uncertainty of orders that are larger than expected (Bartezzaghi et al., 1999b). ARIMA models typically require a long history of observations over time. They were popularized by Box and Jenkins (1970) and have become a standard method in forecasting software. These methods are often applied as if the demand of each product is independent of other products.

#### Significance of the Study

Inventory control for fast-moving products has been studied extensively (Hollier et al., 2002). In sharp contrast, slow-moving inventory has not received equal attention due in part to the lack of observable historical sales figures. Slow-moving products are often overlooked because they do not provide the bulk of sales although they often make up the bulk of the inventory on hand.

With global competition becoming more intense, companies today cannot grow at any cost. Focusing on efficient management of high revenue generating products produces short-term results. Companies may easily carry inappropriate quantities of slow-moving products in which future demand rates are difficult to forecast. Because of the importance of maintaining a given service level as suggested by Miragliotta and Staudacher (2004), organizations can compensate for poor forecasts by increasing assets or working capital, but these options may prove costly. Being able to accurately predict demand for slow-moving products

and maintain an optimal inventory level of them is one way to cut waste.

Unfortunately, traditional forecasting techniques often result in stocking higher than needed levels of inventory for slow-moving products.

Improved methodologies for constructing reliable estimates of future demands of slow-moving retail products are desirable for several reasons. First, sometimes products, such as belts at a men's clothing store, need to be carried because they are expected to be available. A company may even have a critical product that must be maintained although the sales are sporadic. In these cases, management must determine an appropriate level of inventory to carry. Second, the underlying demand distribution may be difficult to estimate if observable sales are not available. More efficient techniques for predicting future demands for slow-moving inventories can lower holding costs, minimize obsolescence, reduce required working capital, increase cash flows, as well as improve the ability to fulfill customer orders.

This dissertation provides information about the performance of proposed inventory managements tools and the conditions under which these tools are reliable and useful in inventory decision making concerning products having low demand rates. The dissertation addresses the situation in which a manager may need an estimate of the future demand rate of products that have not sold over a specified time frame. In addition, an approach that updates estimates of demand rates using an assumed prior distribution on demand rates of products is examined under cost ratios for shortages and surpluses of inventory. An implicit

assumption made in this dissertation is that demand for each product is independent of the demand for other products.

The remainder of this dissertation is arranged as follows. Chapter 2 conducts a comprehensive literature review of relevant literature. Chapter 3 examines the major methodologies employed and narrows the scope of the research. Chapter 4 presents the key findings of the study. Finally, Chapter 5 draws the conclusions, identifies the managerial implications, and describes areas for further investigation.

## CHAPTER 2

### LITERATURE REVIEW

#### Inventory Models

Existing inventory models are based on a variety of assumptions such as variability of demand, cost structures, and physical characteristics of the system (Lee & Nahmias, 1993). They are often classified into two groups based on the nature of demand: deterministic and stochastic (Haddock et al., 1994). For inventory management purposes, deterministic or steady state demand is easier for retail managers to respond to than stochastic demand, which unfortunately, is what most retailers experience.

While somewhat restrictive, deterministic inventory models often provide a good starting point for solving more complicated inventory problems. Nevertheless, the stochastic models and their underlying assumptions are of primary interest to researchers in this area. When demand is constant, the state of the system can be determined at any time for a known order quantity and a reorder point. However, when demand is stochastic, accurate predictions are not possible since the times of occurrence are random variables and simplifying assumptions must be made about the demand pattern to make the problem tractable (Haddock et al., 1994).

Many modern inventory management and control software packages assume that future demand follows a normal distribution and thus include analytic models believed to have sufficient forecasting accuracy (Vereecke &

Verstraeten, 1994). However, that assumption does not hold when the inventory moves very slowly, making it difficult to determine the demand rate.

### Periodic and Continuous Replenishment Models

Stochastic mathematical models form the basis of many inventory systems to address such issues as optimal time for order replenishment and optimal order size to minimize overall costs (Lee & Nahmias, 1993). Most stochastic inventory models use either periodic or continuous review (Stevenson, 2007). The periodic review approach determines the amount of inventory at intermittent intervals to compute the order quantity required to bring it back to a desired level, which is popular with small retailers. The continuous system approach, in contrast, updates inventory instantly and submits an order for a fixed quantity to minimize the total cost whenever it reaches a predetermined minimum level.

Lee and Nahmias (1993) describe variations of single product, single location inventory models but note that the assumptions made about them are mainly related to demand behaviors, costs, and physical characteristics of the system. Generally, the distribution of demand over time is the most important factor in formulating an optimal order policy. It is also necessary to take into account average or discount holding costs, ordering costs, cost changes, and penalty costs. Other key considerations include lead-time, backordering policies, review process, and effect of storage.

## Continuous Review: (s, S) Model and (s, Q) Model

Under a continuous review inventory system for a single item with stochastic demand, either the order point, order-up-to or (s, S) model or the order point, order quantity or (s, Q) model may be used (Schultz, 1989). In other words, whenever the inventory level drops to s units, an order is released to either bring inventory up to S units or acquire Q units to replenish the stock. While rare in practice, a basic assumption of the (s, S) model is that the demand and relevant cost parameters, do not vary over time. Cohen, Kleindorfer, Lee, and Pyke (1992) present optimal policies in an (s, S) model for stocking spare parts and part families in a multi-echelon distribution system.

Despite being difficult to apply, (s, Q) models are common in the existing literature and often recommended for spare parts inventory and high value, low volume items (Razi & Tarn, 2003). Silver (1991) introduces a simple graphical tool that is appropriate for slow-moving items in an (s, Q) model and can be used to reduce inventory control costs. Chang et al. (2001) develop an (s, Q) model for slow-moving inventory under the assumption of Laplace demand.

### Assumptions of (s, S) Model

Several algorithms have been proposed to determine the optimal (s, S) policy under different scenarios (Wagner, O'Hagan, & Lundh, 1965; Johnson, 1968; Schneider, 1978; Ehrhardt, 1979; Freeland & Porteus, 1980). In general, the following four standard assumptions are made.

1. Demand is stationary for some period of time.
2. Demand is discrete.

3. Demand can be back-ordered.
4. Optimization occurs when the average cost is minimized.

Other assumptions include a fixed lead-time and a compound Poisson process (Archibald & Silver, 1978).

#### Assumptions of (s, Q) Model

Watson (1987) identifies the following common assumptions for the (s, Q) policies.

1. Lead-time varies with a known probability distribution.
2. The standard deviation of lead-time demand is small compared to Q.
3. Orders are received as they are placed.

Although these assumptions are valid for most inventories, they are inappropriate for slow-moving products with varying demands (Watson, 1987). Further research by Chung and Hou (2003) and Chung and Huang (2003) investigates the assumptions of the (s, Q) policies in relation to slow-moving inventory. These articles conclude that the Laplace distribution is appropriate for lead-time demand when dealing with slow-moving inventory and provide a methodology for determining safety stock levels to minimize costs making the (s, Q) model appropriate for slow-moving items under some conditions.

#### Assumptions of (S - 1, S) Model

The one-for-one replenishment or (S - 1, S) model is a special case of the (s, S) model with  $s = S - 1$ , where an order is placed to return the inventory level to S any time a customer demand occurs (Lee & Nahmias, 1993). The terms base stock and order-for-order are also used to refer to such an inventory

system. A more general form of the model uses a variable replenishment quantity,  $y$ , and is denoted by  $(S - y, S)$ . It follows that the  $(s, S)$  and the  $(s, Q)$  models are identical when the demand transactions are unit sized making  $Q$  equal to  $S - y$  (Razi & Tarn, 2003). Gelders and Groenweghe (1985) suggest the  $(s - 1, S)$  model is appropriate for expensive slow-moving items and they provide a detailed description of it along with the underlying assumptions.

### Inventory Models for Spare Parts

A special class of slow-moving inventory is spare parts. Research papers including Gelders and Van Looy (1978), Bartakke (1981), Popovic (1987), Vereecke and Verstraeten (1994), Weingart (1991), Klein Haneveld and Teunter (1997), Dekker, Kleijn, and Rooij (1998), Strijbosch, Heuts, and Schoot (2000), Razi and Tarn (2003), and Dolgui and Pashkevich (2007) propose models for managing spare parts, which is important in ensuring adequate availability of items, thus providing organizations with a competitive advantage. (Additional papers are included in Figure 4.) Many of the analytical models also apply to slow-moving inventory (Cohen & Lee, 1990).

Gelders and Van Looy (1978) apply the Pareto Principle to classify spare parts and determine optimal inventory policies for fast movers, normal movers and slow movers. Bartakke (1981) develops a mathematical model and a simulation technique based on a cost function for an  $(s, S)$  model for groups of spare parts. Klein Haneveld and Teunter (1997) provide optimal strategies for stocking spare parts for expensive equipment. Dekker et al. (1998) develop separate models for critical and non-critical parts when demand follows a

Poisson distribution and the stocking policy is lot-for-lot, which is when demand occurs a replacement part is ordered.

Research on the topic of inventory management of spare parts includes studying competing models, forecasting methods, control procedures, and demand distributions. A summary of the literature on spare parts inventories is presented by Kennedy, Wayne Patterson, and Fredendall (2002). When demand is uncertain for consecutive periods, Popovic (1987) discusses a periodic replenishment model for spare parts where appropriate parameters are adjusted as periods of demand accumulate. Vereecke and Verstraeten (1994) study spare parts in a chemical plant, propose the Poisson distribution for modeling demand and design an algorithm for effective inventory control. Strijbosch et al. (2000) introduce a forecast-inventory control procedure to model lumpy and sporadic demands for spare parts at a MARS candy production plant in the Netherlands.

Ghobbar and Friend (2002), Ghobbar and Friend (2003), and Ghobbar (2004) compare thirteen forecasting methods for spare part demand for airline fleets and find that the overall best approach is Croston's (1972) weighted moving average technique (see the section of Demand for Slow-Moving Inventory below for details). Eaves and Kingsman (2004) also examine several forecasting methodologies for slow-moving military spare parts and demonstrate that a modification of Croston's method lead to a reduction of inventory.

Willemain et al. (2004) compare several approaches to forecast demand for spare parts in manufacturing firms: Bootstrap method, Croston's method, and exponential smoothing.

Table 4

## Research on Spare Parts Inventory

Author(s)	Year	Bayesian Model	Proposed Model	Model Comparison	Industry
Smith and Vemuganti	1969	X			None
Haber and Sitgreaves	1970		X		Military-Naval
Brown and Rogers	1973	X			F-14 Aircraft
Burton and Jaquette	1973	X			None
Gelders and Van Looy	1978		X		Petrochemical Plant
Bartakke	1981		X		Computer Hardware
Popovic	1987	X			None
Weingart	1991		X		Military-Naval
Cohen et al.	1992		X		None
Vereecke and Verstraeten	1994		X		Chemical Plant
Klein Haneveld and Teunter	1997		X		None
Sani and Kingsman	1997			X	Agricultural Machinery
Dekker et al	1998		X		Petrochemical Plant
Strijbosch et al.	2000				Manufacturer
Ghobbar and Friend	2002			X	Airline
Kennedy et al.	2002			X	Overview
Ghobbar and Friend	2003			X	Airline
Razi and Tarn	2003		X		Manufacturer

Table 4 (continued)

Research on Spare Parts Inventory

Author(s)	Year	Bayesian Model	Proposed Model	Model Comparis	Industry
Ghobbar	2004			X	Airline
Eaves and Kingsman	2004				Military
Willemain et al.	2004		X		Various Industries
Syntetos and Boylan	2005	X	X		Automotive
Syntetos and Boylan	2006	X	X		Automotive
Dolgui and Pashkevich	2007		X		None

Syntetos and Boylan (2006) empirically compare simple moving average, single exponential smoothing, Croston's method, and a modified Croston method proposed by Syntetos and Boylan (2005) on spare parts from the automotive industry with a focus on intermittent demands. Syntetos and Boylan's (2006) proposed methodology reduced inventory costs on average compared to the other methods. The techniques discussed in this section are summarized in Figure 4 and categorized by their approach to examining models.

#### Inventory Model Cost Structures

Thompstone and Silver (1975) consider the following cost structure for keeping an item in stock:  $\text{Expected Cost} = \text{Setup Cost} + \text{Carrying Cost}$ , where setup or ordering cost is a fixed amount and carrying or holding cost is normally a

percentage of the product's unit cost. Razi and Tarn (2003) propose a cost model with the expected annual inventory cost equal to the review cost per year plus the ordering cost per year and the annual carrying cost. A direct search method is recommended for finding the optimal  $(s, S)$  policy. Their suggested approach was demonstrated to be appropriate for slow-moving items.

Popovic (1987) determines the optimal stock level per unit time interval by specifying the unit surplus cost and the unit shortage cost. This paper makes several assumptions including that ordering is done at the beginning of the period, lead-time is zero, and demand can be back-ordered. Williams (1983) provides tables of computed stock-out costs for items with periodic demand when shortage costs are not known.

#### Demand for Slow-Moving Inventory

Research on slow-moving inventory items has been around for a long time and its roots date back to the works of Whitin and Youngs (1955) and Heyvaert and Hurt (1956). The seminal paper by Croston (1974) laid the foundation for further studies in this area. Heyvaert and Hurt (1956) determine an optimal inventory level minimizing the cost of storage and the cost of customer dissatisfaction. However, their method is restricted to slow-moving items with a Poisson demand pattern and long replenishment periods.

Feeney and Sherbrooke (1966) derive analytical solutions for inventory demand with a compound Poisson distribution and a reorder quantity of one. Building on the findings of Feeney and Sherbrooke (1966), Croston (1974) addresses the key issue of whether or not to stock a slow-moving product and he

proposes separate forecasts for demand size and inter-demand intervals to overcome the biases caused by periods of no demand. More recently, Shenstone and Hyndman (2005) have proposed stochastic models underlying Croston's method for slow-moving items and they show the point forecasts and prediction intervals based on such models to be useful.

### Forecasting Methods

Predicting demand for slow-moving inventory is a major challenge since historical data are often limited. Unfortunately, many approaches to increasing productivity, such as MRP II, depend on accurate forecasts (Willemain et al., 1994). Traditional methods for solving this problem are proposed by Haber and Sitgreaves (1970), Burton and Jaquette (1973), Croston (1972), and Williams (1984). They generally suggest different ways of classifying products and estimate demand for items in each category by using an appropriate technique.

Silver (1965), Brown and Rogers (1973), and Smith and Vemuganti (1969) offer Bayesian approaches to forecasting the demand for slow-moving goods based on historical data, which are often not available. Muckstadt and Thomas (1980), Silver (1970), and Thompstone and Silver (1975) have suggested additional models for special cases involving slow-moving items. Levén and Segerstedt (2004) propose an ERP system implementation based on the Croston method that is capable of forecasting slow-moving and fast-moving products, however their approach was criticized for being biased by Boylan and Syntetos (2007). In addition to emphasizing the importance of the distribution of demand,

Sani and Kingsman (1997) also suggest that the variance of the forecast is an important factor in determining a proper method.

Popovic (1987) delineates inventory models under the following assumptions.

1. Demand rate is unknown but constant and demand follows a Poisson distribution.
2. Demand rate is unknown but variable.
3. Demand rate is unknown and demand follows a distribution from a family of distributions.

Willemain et al. (2004) develop a distribution independent bootstrap method for estimating intermittent demand, which was demonstrated to outperform some of the above-mentioned methodologies.

Table 5 provides a summary of articles related to demand forecasting methodologies discussed in this section and the dominant methodology contribution of each paper. The bootstrap method, recently demonstrated to be a competitive procedure, did not receive much attention until recent years as is noted in Table 5. In contrast, Bayesian approaches have appeared over a long span of research in forecasting demand for products.

Table 5

## Demand Forecasting Methodologies

Author(s)	Silver	Smith and Vemuganti	Haber and Sitgreaves	Silver	Croston	Brown and Rogers	Burton and Jaquette	Thompstone and Silver	Muckstadt and Thomas	Williams	Popovic	Sani and Kingsman	Leven and Segerstedt	Willemain et al.
Year	1965	1999	1999	1999	1999	1999	1999	1999	1999	1999	1999	1999	2000	2000
Derivation of demand from														
similar products			X				X			X				
Bayesian approach	X	X				X					X			
Bootstrap method														X
Others				X	X			X	X			X	X	

## Lumpy Demand

Lumpy demand is characterized by sporadic or intermittent requirements over time. The existence of lumpy demand patterns in slow-moving inventory has been documented by several of researchers (Croston, 1972; Ward, 1978; Schultz, 1987; Vereecke & Verstraeten, 1994; Strijbosch et al., 2000). Inventory with a smooth and continuous demand has been thoroughly investigated and a large number of models have been proposed; however, many of them do not

apply to products with lumpy demand (Williams, 1984; Ward, 1978; Schultz, 1987; Willemain et al., 1994).

Lumpy demand is extremely irregular and highly variable as characterized by a large coefficient of variation as well as sales peaks followed by periods of no or low sales (Bartezzaghi et al., 1999b). Zotteri (2000) discusses factors that contribute to the lumpiness. Bartezzaghi, Veganti, and Zotteri (1999a) provide a methodology for dealing with lumpy demand based on the sources of lumpiness, such as number of potential customers, heterogeneity of customers, customer buying patterns, and correlation between customer requests.

Sophisticated ERP systems that accurately forecast the demand of fast-selling merchandise generally rely on the exponential smoothing technique but they do not perform as well for slow-moving items (Willemain et al., 1994; Syntetos & Boylan, 2001; Razi & Tarn, 2003). This is because exponential smoothing places more weight on current data. As such in the case of lumpy demand estimates based on periods of no demand right before a period with a positive demand will be biased (Syntetos & Boylan, 2001).

Croston (1972) demonstrates that for intermittent-demand items the forecast error may be reduced by smoothing the times between nonzero demands and the demand sizes separately. His method remains relevant and practical. However, despite its theoretical superiority, Syntetos and Boylan (2001) point out actual improvement in performance was modest or lacking when the forecasting technique was adopted by and implemented in some organizations.

The study of lumpy demand is a niche area of operations management and has emerged from research on slow-moving items (Miragliotta & Staudacher, 2004). Lumpy patterns are exhibited by stock that has little or no sales activity for several periods and the demand typically cannot be accurately forecasted by standard inventory control (Silver, 1970). Wilcox (1970) notes that many techniques minimizing the mean absolute deviation provided unsatisfactory results with lumpy demand items and he suggests a simple method based on the desired service level to determine order quantities.

Ghobbar and Friend (2002) demonstrate that the sources of lumpiness must be identified and understood for accurate predictions of lumpy demand for airline service parts. Regattieri, Gamberi, Gamberini, and Manzini (2005) apply 20 approaches to lumpy demand series of data and find that Croston's method and exponentially weighted moving average to be superior for forecasting lumpy demand.

While the research streams on slow-moving demand and lumpy demand share similar roots and some common problems, lumpy demand has extreme fluctuations in variance in addition to fluctuations in mean demand rate (Miragliotta & Staudacher, 2004). The issue of determining the values of inventory policy parameters has been addressed by a few researchers on lumpy demand (Williams, 1982; Ward, 1978; Archibald & Silver, 1978). While most of them are variations of the classic  $(s, S)$  model, Watson (1987) argues that several of the simplifying assumptions are violated by the irregularity of slow-moving inventory demand patterns. Additional methods for forecasting

intermittent demand are developed by Williams (1982), Schultz (1987), Watson (1987), and Dunsmuir and Snyder (1989).

#### Most Recent Research Related to Slow-Moving Inventory

A stream of research (Syntetos & Boylan, 2005; Shale, Boylan, & Johnston, 2006; Syntetos, Boylan, & Croston, 2005) related to slow-moving inventory have compared competing forecasting methods, proposed modifications of proposed methods, and examined inventory methods for the management of spare parts. Syntetos and Boylan (2005) compare simple moving averages, single exponential smoothing, Croston's method and a proposed modified Croston's method, which dominated the comparisons. Shale et al. (2006) propose an adjustment method for using exponentially weighted moving average to forecast demand for slow-moving, intermittent demand.

Hua, Zhang, Yang, and Tan (2007) propose using explanatory variables to forecast slow-moving spare parts and compare these forecasts to those from exponential smoothing, Croston's method, and bootstrapping. In order to adopt an appropriate method for forecasting slow-moving inventory, Syntetos et al. (2005) propose using the average mean time between demand and the squared coefficient of variation of demand size to categorize demand. They suggest that determining the category of demand is the first step for adopting the most appropriate forecasting methodology.

#### Demand Distributions

Although both pioneered by Whitin and Youngs (1955) as well as Heyvaert and Hurt (1956), research on lumpy demand and on slow-moving

demand have progressed in different directions (Williams, 1984). The fundamental problem with slow-moving products, including spare parts and maintenance items, is that little data are available even when sales occur (Haber & Sitgreaves, 1970). As described in the previous section, lumpy demand is more volatile than slow-moving demand and many probability distributions have been proposed for modeling them, including Poisson, exponential, lognormal, gamma, logistic, negative binomial, and Pearson (Silver, Pyke, & Peterson, 1998). Figure 6 provides a summary of related research and the demand distributions employed. While the normal distribution is often assumed for demand of most items in inventory systems, the Poisson distribution is often used to model the demand for slow-moving items (Vereecke & Verstraeten, 1994; Bagchi et al., 1983). The next few sections discuss several of these distributions.

#### Poisson Distribution of Demand for Slow-Moving Inventory

Heyvaert and Hurt (1956), Hadley and Whitin (1963), Gelders and Van Looy (1978), Schultz (1987), and Silver et al. (1998) recommend the Poisson distribution for modeling the demand patterns for slow-moving inventory. This distribution is generally appropriate provided that the demand variance falls within 10 percent of the mean (Silver et al., 1998). However, Vereecke and Verstraeten (1994) remark that use of the Poisson distribution to model business data often violates this condition. Silver (1970) describes interesting applications of the Poisson distribution within military and industrial organizations and finds it to be most useful when both the demand and the number of orders are large.

Since the Poisson process is germane to so much of the methodology studied about slow-moving inventory and inventory management, in general, the definition of a Poisson process follows. A Poisson process is a set of random variables indexed over time. Each random variable  $X$  has a Poisson distribution as described by (1).

$$P\{X = x\} = (\lambda t)^x \frac{e^{-\lambda t}}{x!}, \quad t \geq 0, \quad x = 0, 1, 2, \dots \quad (1)$$

Ross (2002) defines a Poisson process as a counting process  $\{N(t), t \geq 0\}$  with a positive rate  $\lambda$ , if the next three conditions hold.

1.  $N(0) = 0$ .
2. The process has independent increments.
3. The number of events in any interval of length  $t$  is Poisson distributed with mean  $\lambda t$ . In addition, for all  $s, t \geq 0$ ,

$$P\{N(t+s) - N(s) = n\} = e^{-\lambda t} \frac{(\lambda t)^n}{n!}, \quad n = 0, 1, 2, \dots \quad (2)$$

Furthermore,  $E[N(t)] = \lambda t$ .

#### Compound Poisson Distribution of Demand for Slow-Moving Inventory

Feeney and Sherbrooke (1966), Thompstone and Silver (1975), and Archibald and Silver (1978) suggest that the compound Poisson distribution be considered in modeling customer demand for slow-moving items. A good example of this distribution is a scenario where customers arrive at a store according to a Poisson process, each of them purchases batches of products and the total number of product units bought follows a compound Poisson distribution. Vareecke and Verstraeten (1974) propose a forecasting procedure

based on the compound Poisson distribution by assuming that each occurrence of demand consists of a package of goods and each package contains multiple units.

Feeney and Sherbrooke (1966) explain that the compound Poisson distribution can be visualized in an inventory problem as a series of customers with Poisson arrivals who demand an amount  $D$  that has an independent discrete distribution with probabilities  $P\{D = j\}$  where  $j = 0, 1, \dots$  is the amount demanded. The probability that  $y$  customers order a total of  $x$  demands is the  $y$  fold convolution of the distribution with probabilities  $P\{D = j\}$  which is denoted by  $f_y(x)$ . By convention and without any loss of generality,  $P\{D = 0\} = 0$ . Using this notation, the compound Poisson probability of  $x$  demands is presented in (3) with  $\lambda$  representing the customer arrival rate. The summation index  $y$  is incremented over the number of customers that order  $x$  demands and thus,  $0 \leq y \leq x$ .

$$P\{X = x | \lambda\} = \sum_{y=0}^x \left( \frac{\lambda^y e^{-\lambda}}{y!} \right) f_y(x), \quad 0 \leq \lambda < \infty, \quad 0 \leq x < \infty \quad (3)$$

A special case of the compound Poisson distribution is the stuttering Poisson distribution (Feeney & Sherbrooke, 1966), which has been used to model customer buying patterns when demand for a product is lumpy (Hollier et al., 2002). Ward (1978) applies the stuttering Poisson distribution to calculate inventory reorder points in situations where the sporadic demand follows a Poisson distribution but the order size follows a geometric distribution.

## Gamma Distribution of Demand for Slow-Moving Inventory

Several authors have used the gamma distribution to model the sales of slow-moving items (Burgin, 1975; Snyder, 1984; Dunsmuir & Snyder, 1989; Segerstedt, 1994; Tyworth et al., 1996; Yeh, 1997). The gamma distribution has some interesting properties, and one of them is that a change in the mean may affect the variance. Moreover, both the exponential distribution and the chi-square distributions are special cases of the gamma distribution. If the random variable  $X$  is gamma distributed, it has the following probability density function, with mean and variance equal to  $\frac{\alpha}{\beta}$  and  $\frac{\alpha}{\beta^2}$ , respectively.

$$f(x; \alpha, \beta) = \begin{cases} \frac{\beta^\alpha x^{\alpha-1} e^{-\beta x}}{\Gamma(\alpha)}, & 0 < x < \infty \\ 0 & , \text{ elsewhere} \end{cases} \quad (4)$$

$$\alpha > 0, \beta > 0$$

Segerstedt (1994) justifies the applicability of the gamma distribution to modeling the movement of merchandise and explains that the distribution is generally mathematically tractable in its inventory control applications. Furthermore, he argues that the gamma distribution is a better fit to the distribution of sales data than the normal distribution, particularly since the time between sales or the demand size cannot be negative.

Findings from the research of Tyworth et al. (1996) provide inventory planners with an optimum reorder point and a lot size that minimizes the total costs using a nonlinear programming approach for the  $(s, Q)$  model under

continuous review. Their study considered service levels and random lead times in recommending the gamma distribution for modeling the demand for A, B, or C type of inventory items.

Dunsmuir and Snyder (1989) in a case study identify lumpy demand patterns in spare parts inventory. When products were pooled by type, the authors decided that the gamma distribution was more appropriate to model demand than the normal and Poisson distributions. Yeh (1997) chooses the gamma distribution over the Poisson distribution to represent the demand for spare parts in a mid-sized Taiwanese electronic company due to the large variations in their demand rates.

#### Normal Distribution of Demand for Slow-Moving Inventory

Croston (1974) studies the basic decision to stock or not stock slow-moving items with the aid of a cost formula. This work is based on the assumptions that demand follows a normal distribution and that replenishment occurs at fixed intervals in order to maintain a pre-specified inventory level.

The mound-shaped or approximately mound-shaped probability distributions of business data are often modeled using the normal distribution. When inventory levels are subject to continuous review, demand is conveniently assumed to follow a normal distribution. Ross (2002) defines the following probability density function for the normal random variable  $X$  with mean  $\mu$  and variance  $\sigma^2$  as:

$$f(x; \mu, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-\mu)^2}{2\sigma^2}}, -\infty < x < \infty, -\infty < \mu < \infty, \sigma^2 > 0 \quad (5)$$

Despite a litany of limitations, Croston (1972), Bartakke (1981), Vereecke and Verstraeten (1994), and Burgin (1972) use the normal distribution to model the demand pattern for slow-moving inventory. Since this distribution assumption allows negative values, it is only applicable when the coefficient of variation of the lead-time demand is 0.2 or less (Snyder, 1984). With slow-moving products having zero demand for an extended period of time, the normality assumption often causes difficulties in model interpretation (Vereecke & Verstraeten, 1994).

#### Hermite Distribution of Demand for Slow-Moving Inventory

Bagchi et al. (1983) introduce the use of the Poisson-like Hermite distribution to model Poisson-distributed demand during normally distributed lead time for slow-moving items. They maintain the Hermite distribution is as simple as the Poisson distribution. Its probability density function is as follows with  $p_w$  representing the probability that a Hermite variable assumes the value of  $w$ :

$$p_0 = e^{-(a+b)},$$

$$p_w = p_0 \sum_{j=0}^{\left[\frac{w}{2}\right]} \frac{a^{w-2j} b^j}{(w-2j)! j!}, \quad w = 1, 2, \dots \quad (6)$$

where  $a = \lambda\mu - \lambda^2\sigma^2$ ,  $b = \frac{\lambda^2\sigma^2}{2}$ ,  $w$  is the demand during the lead time with mean

$\lambda$ ,  $\mu$  is the mean of the lead time, and  $\sigma^2$  is the variance of the lead time. Note

that when  $\frac{w}{2}$  is non-integral, the index  $j$  is incremented to the largest integer less

than or equal to  $\frac{w}{2}$ . Furthermore, the distribution is valid provided  $a \geq 0$  or

$\frac{\mu}{\sigma^2} \geq \lambda$ . Bagchi et al. (1983) also provide a simple formula for computing the mean and standard deviation of a Hermite distribution.

#### Other Distributions of Demand for Slow-Moving Inventory

Snyder (1984) successfully applies the economic order quantity (EOQ) approach to slow-moving inventory with a demand following the exponential distribution. Janssen, Heuts, and Kok (1998) discuss the (R, s, Q) inventory model when known intermittent demand is modeled as a compound Bernoulli process. They present an approximation method to compute the reorder point since neither R nor Q is mentioned with a service level restriction.

An inventory management system for slow-moving and fast-moving products is proposed by Levén and Segerstedt (2004). Their model is based on the Croston method and sporadic demand is assumed to have an Erlang distribution. Silver et al. (1998) provide a detailed description of the exponential, logistic, lognormal, negative binomial, and other distributions and offer insights into the appropriate use of these distributions. A summary of the research discussed in this section is shown in Figure 6. The research articles in this figure are classified by the primary demand distributions that were investigated.

Table 6  
Demand Distributions Cited in Literature

Author(s)	Year	Poisson	Compound Poisson	Gamma	Normal	Hermite	Exponential	Logistic	Negative binomial	Erlang
Heyvaert and Hurt	1956	x								
Hadley and Whitin	1963	x								
Feeney and Sherbrooke	1966		x							
Burgin	1972				x					
Croston	1977				x					
Burgin	1977									
Thompson and Silver	1977		x							
Archibald and Silver	1977		x							
Ward	1977		x							
Gelders and Van Looy	1977	x								
Bartakke	1981					x				
Bagchi, Hayya, and Ord	1983						x			
Snyder	1984							x		
Schultz	1987	x								
Dunsmuir and Snyder	1988									
Segerstedt	1989									
Vereecke and Verstraeten	1994									
Tyworth, Guo, and Ganeshan	1999									
Yeh	1997									
Silver, Pyke, and Peterson	1999	x								
Hollier et al.	2002									
Levén and Segerstedt	2002									x

### Bayesian Approach to Inventory Modeling

Popovic (1987) proposes a Bayesian approach to inventory decision making, which allows the estimates of the parameters of the a priori distribution of demand rate  $\lambda$  to be updated. For example, if two time periods of recording sales have passed, then the a posteriori distribution of  $\lambda$  is gamma with parameters  $\alpha > 0$  and  $\beta > 0$  given that values of the demands  $X_1$  and  $X_2$ , namely,

$\lambda | X_1, X_2 \sim \Gamma(\alpha + X_1 + X_2, \beta + 2)$ . Popovic (1987) states that the optimal inventory levels should be determined by using the a posteriori distribution and knowledge of the surplus cost per unit of time,  $C_1$ , as well as the shortage cost per unit of time,  $C_2$ .

Silver (1965) applies the Bayesian method to select the reorder point for an inventory model. De Wit (1983) proposes a Bayesian approach to forecasting slow-moving items; however, slow moving was defined as 10 or fewer demands per unit of time, higher than the demand rate used in this research. Furthermore, the proposed method failed to perform well when demand was extremely low. Price and Haynsworth (1986) suggest that the Bayesian approach is better suited to predict the sales of products with slow demand than exponential smoothing although its actual performance may depend on the distribution of the demand. A case study is presented by Aronis, Magou, Dekker, and Tagaras (2004) utilizing a Bayesian approach for spare parts, however the research was not specifically focused on slow-moving demand.

### Software Reliability

The prediction of software reliability uses some methodologies that maybe adapted for the study of slow-moving inventory. Many models based on the Poisson distribution have been considered for estimating failure rates in software (Abdel-ghaly, Chan, & Littlewood, 1986; Kaufman, 1996). As computer software packages are developed, rigorous testing is required to identify faults or bugs before they are released to the customer. The rates of the occurrence of bugs have been found to follow a variety of distributions in addition to the Poisson,

including the exponential, gamma, Weibull, and geometric distributions (Abdel-ghaly et al., 1986; Miller, 1986). However, Abdel-ghaly et al. (1986) suggest that no one distribution has been shown to be superior to others in forecasting the number of bugs a priori, leaving the programmer to determine which distribution best fits the historical data according to its predictive ability.

The inherent complexity of the software development process is created by many factors. As such, estimating the reliability of individual software packages or those on a client's server network may be an elusive task even after design review, module testing, and self-checking. Software engineers typically put the software through a series of tests phase to remove any bugs found and determine when it is ready for use. However, in the interest of marketing the product in a timely fashion, the software is generally released without complete knowledge of the hidden faults (Lyu, 1995).

The term bug is defined by Omdahl (1988) as a program defect, which is used here as an equivalent for fault. According to him, a fault is a defective, missing, or extra instruction or set of related instructions that is the cause of one or more actual or potential failure types. A fault with the software does not necessarily render the system inoperable.

There is intense pressure on software companies to introduce new products to the market even knowing that some bugs (for example, errors in coding) are present. They often attempt to estimate the rate of faults occurring in the software without having observed the errors made by the remaining bugs. Towards that end, assumptions about the distribution of the number of errors

need to be made. Schulmeyer and McManus (1999) note that several distributions with exponentially decreasing probability density functions have been used in software reliability research, including logistic, hypergeometric, and Rayleigh. Additionally, Ross (2002) proposes a statistical method for estimating software reliability based on the occurrence of bugs following a Poisson distribution. A detailed account of Ross' (2002) approach will be given in the remainder of this section.

Suppose that there are  $n$  bugs contained in a software package. The number of errors caused by bug  $i$  is assumed to follow a Poisson distribution with a mean of  $\lambda_i$ ,  $i = 1, 2, \dots, n$ . Ross (2002) defines  $\Psi_i(t) = 1$  if bug  $i$  has not caused a detected error by time  $t > 0$  and 0 otherwise,  $i = 1, 2, \dots, n$ . These indicator variables allow the remaining error rate to be expressed as  $\Lambda(t) = \sum_{i=1}^n \lambda_i \Psi_i(t)$ .

Estimating this quantity is important in the decision-making process of releasing the software to the market. Obviously, a high error rate would be unacceptable to the customer.

To estimate  $\Lambda(t)$ , Ross (2002) uses  $M_j(t)$  to denote the number of bugs that are responsible for  $j$  detected errors by time  $t$ ,  $j = 1, 2, \dots, n$ . For instance  $M_1(t)$  is the number of bugs that cause exactly one error,  $M_2(t)$  is the number of bugs that cause exactly two errors, and so on with  $\sum_{j=1}^n j M_j(t)$  being the total number of detected errors. It can be shown that  $E[\Lambda(t) - \frac{M_1(t)}{t}] = 0$ , which is key to establishing that  $\frac{M_1(t)}{t}$  is an unbiased estimate of  $\Lambda(t)$ . As will be

demonstrated in Chapter 3, the second moment of  $\Lambda(t) - \frac{M_1(t)}{t}$  is the same as

the expected value of  $\frac{M_1(t) + 2M_2(t)}{t^2}$ , which is a function of  $M_1(t)$  and  $M_2(t)$ .

Therefore, the mean squared difference between  $\Lambda(t)$  and  $\frac{M_1(t)}{t}$  can be

estimated by  $\frac{M_1(t) + 2M_2(t)}{t^2}$ . In the context of estimating demand for slow-

moving inventory, the notation mentioned for software reliability can be interpreted as shown.

$n$  = Number of products

$\lambda_i$  = Demand rate for product  $i$ ,  $i = 1, 2, \dots, n$

$t$  = Length of time period over which demand is observed

$\Psi_i(t)$  = Equal to 1 if product  $i$  has no sales by time  $t$  and 0 otherwise

$\Lambda(t)$  = Theoretical demand rate of products experiencing no sales

$M_j(t)$  = Number of products that have sold exactly  $j$  items by time  $t$

$\frac{M_1(t)}{t}$  = Estimator of  $\Lambda(t)$

$\frac{M_1(t) + 2M_2(t)}{t^2}$  = Estimator of mean squared difference between  $\frac{M_1(t)}{t}$  and  $\Lambda(t)$

In this study,  $\Lambda(t)$  will represent the demand rate for slow-moving products that have not experienced a previous sale over a specified period of time. By definition,  $\Lambda(t)$  is the sum of the theoretical underlying demand rates for unsold products. Note that this procedure is easy to implement since  $M_1(t)$  and  $M_2(t)$  are simply counting functions. The relationship between Ross' (2002) methodology

for software reliability and slow-moving inventory applications is discussed in Chapter 3.

Ross (1985a) makes two basic assumptions: bugs are independent and there is a probability  $p$  of detecting a bug. In the context of slow-moving inventory, these two assumptions are incorporated by assuming that the demand of each product is independent of the demand of other products and that the probability of detecting a sale is equal to one, that is, sales are detected with 100% certainty. Prediction intervals for the future demand rate of slow-moving inventory will be proposed and assessed using these concepts.

#### Decision Making Involving Slow-Moving Inventory

As companies are striving to operate efficiently in today's global market, more research is needed to help managers accurately forecast future demand, so that they will be able to make sound decisions on slow-moving inventory that experiences independent demand occurrences. Johnston et al. (2003) investigate actual orders with an electrical wholesaler to study the effect of order size from customers. Rosenfield (1989) addresses the problem of disposing of slow-moving or obsolete inventory. Teunter and Vlachos (2002) study the effect of low demand on remanufacturing. Tekin, Gurler, and Berk (2001) examine slow-moving perishable inventory and propose an age-based policy for optimal control. Service parts, which are similar to spare parts, have received special attention from Fortuin and Martin (1999).

Decision rules for obsolescent items are applied to slow-moving inventory by Hummel and Jesse (1990). George (1982) develops methods for discounting

slow-moving inventory to free cash for investing in items that turn over more quickly. Horodowich (1979) introduces an approach for evaluating the decision to write off slow-moving inventory based on the profit loss, the reduction of working capital, and its effect on cash flow.

Recently, several methodologies have been proposed for categorizing slow-moving items since correct identification of demand type can lead to the selection of the most appropriate technique for forecasting demand. For example, Williams (1984) was the first to classifying lumpiness and intermittence based on the squared coefficient of variation of demand size. Eaves and Kingsman (2004) suggest a refinement of Williams' (1984) classification scheme according to the following characteristics of a product's sales behavior: variability of demand, lead time, frequency of sales - slow-moving, smooth, periodic, or erratic. However, neither approach is applicable to fast-moving products.

Johnston and Boylan (1996) propose a categorization system based on the length of inter-demand interval for determining when simple exponential smoothing is better than Croston's method and vice versa. More recently, Syntetos et al. (2005) use the mean square error (MSE) as a criterion to compare different prediction methods and classify demand patterns in terms of the average inter-demand interval and the coefficient of variation. Four categories result from their study: erratic but not very intermittent, lumpy, smooth, and intermittent but not very erratic. Appropriate forecasting techniques have been proposed for each of them.

Given the many research papers that have investigated slow-moving inventory models, a natural question is: What is the unique contribution of the research in this dissertation to this body of knowledge? First, the proposed prediction intervals for future demand rates of products with no sales or no more than one sale per time period and their performance over a variety of conditions has not been previously explored. The literature does not present or assess methodology for addressing the very difficult problem of estimating future demand rates for products without a history of sales.

Second, the Monte Carlo simulation study assessing a Bayesian approach for obtaining optimal inventory levels under conditions of intermittent demand has not been previously presented and is thus a unique contribution to the literature on proposed Bayesian approaches. There are no published simulation studies assessing the performance of these models using slow-moving inventory.

A motivation for the proposed prediction intervals for the demand of slow-moving inventory used in this dissertation is that the approaches used in software reliability may be applicable. In the case of software reliability, the future error rate of a bug must be assessed without any record of previous errors. In the case of inventory with intermittent demand, a manager must assess the future demand rate for products with little or no history of sales over a specified time frame. Although there are many differences between software reliability and inventory management, the proposed intervals are modeled using the same methodology used in software reliability. To determine whether these proposed intervals are

reliable for the inventory management of products having intermittent demand, a Monte Carlo simulation study is conducted and discussed in Chapter 4.

The motivation for assessing a Bayesian approach using a Monte Carlo simulation is that several papers previously cited in this chapter discuss the merits of assuming a prior distribution for the demand rate of products. In particular, the approach presented by Popovic (1987) using a Bayesian approach for a single period and then updating parameters to estimate optimal inventory levels for the next period assuming that surplus and shortage costs are available addresses the problem of stocking spare parts. However, the performance of his approach has not been studied under various conditions and, in particular, not under conditions in which intermittent demand is present.

## CHAPTER 3

### METHODOLOGY

Chapter 2 provided a review of inventory systems and methods for estimating the demand rates of slow-moving products. Characteristics of certain demand distributions were also presented. Next, this background information is applied to the methodologies presented in this dissertation, which are applicable to several situations involving products with low levels of independent demand.

#### Proposed Methodologies

When no demand history is available for slow-moving items, assessing the demand rate at zero will often underestimate the actual demand since it is probably not zero, but a very low value. An evaluation of techniques for estimating the demand rates of slow-moving products is lacking in the literature. This chapter illustrates several proposed methods to fill the void.

The next section describes two procedures for estimating the rate of demand, which is assumed to follow a Poisson distribution. An explanation of the use of these two relatively easy-to-compute formulas is provided. Subsequently, a method is adapted from software reliability literature to provide a better estimate of demand for inventory items with zero sales. To examine the changing relationship between the demand rate and predictions of future demand over time, a Monte Carlo simulation analysis is performed. To provide insights into the distribution of the proposed estimator of future demand rate for products having zero demand over a specified time period, histograms of its simulated values are presented.

A Bayesian approach to estimating future demand rate is also examined. Using this methodology, a simulation study is conducted to compare optimal levels of inventory with respect to the costs for shortages and surpluses to a traditional approach based on the Poisson distribution. The application of this alternative method requires a user-supplied prior distribution for the demand rate, which is treated as a random variable having a gamma distribution.

Stopping rules can be applied to decide on the discontinuation of carrying slow-moving merchandise. If the projected demand rate of a product is not economically high enough to justify carrying the product, for instance, retail managers must act to liquidate its inventory. An easy-to-implement rule is to stop carrying products whose estimated demand rate is below a certain threshold. Alternatively, if the upper endpoint of a one-sided prediction interval for the future demand rate of a group of products is below that threshold, then the decision is to liquidate.

This study assessed the robustness of the reliability of both one-sided and two-sided prediction intervals for the future demand rate of products with no sales history as well as those with no more than one sale. The performance of the Bayesian approach is evaluated using intermittent sales data. The aforementioned methodologies are investigated to answer seven research questions listed at the end of this chapter through a Monte Carlo simulation study. Data from a major national retailer are used to evaluate the proposed approaches.

## Estimating Demand Rate for a Single Product with Nonzero Demand

Since many research studies on optimal levels of inventory often model the product demand as a Poisson process, we will make the same assumption in the remainder of this section. Estimating the demand rate, denoted by  $\lambda$ , appears to be a straightforward exercise and inventory managers may find either of the following procedures to be acceptable.

Estimation Procedure 1: Observe the time  $T$  that it takes for a fixed  $k^*$  number of units to be sold and estimate  $\lambda$  to be  $\frac{k^*}{T}$ .

Estimation Procedure 2: Observe the number of units  $K$  that are sold over a fixed time  $t^*$  and estimate  $\lambda$  to be  $\frac{K}{t^*}$ .

In Estimation Procedure 1, notice that  $k^*$  is fixed and  $T$  is a random variable. Thus, by the assumption of a Poisson process,  $T$  is the sum of  $k^*$  identical exponential random variables distributed as an Erlang( $k^*$ ,  $\lambda$ ), which is a special case of the gamma distribution. Since the likelihood function is proportional to  $\lambda^{k^*} e^{-\lambda t}$ , the maximum likelihood estimator (MLE) is  $\frac{k^*}{T}$ .

For any random variable  $X$  with a gamma distribution,  $\Gamma(\alpha, \beta)$ , in which  $\alpha$  is an integer, the expected value of its reciprocal is presented in (1).

$$\begin{aligned} E\left(\frac{1}{X}\right) &= \int_0^{\infty} \frac{1}{x} \frac{\beta^\alpha x^{\alpha-1} e^{-\beta x}}{\Gamma(\alpha)} dx = \\ \frac{\Gamma(\alpha-1)}{\Gamma(\alpha)} (\beta) \int_0^{\infty} \frac{\beta^{(\alpha-1)-1} x^{(\alpha-1)-1} e^{-\beta x}}{\Gamma(\alpha-1)} dx &= \frac{\beta}{(\alpha-1)} \end{aligned} \quad (1)$$

Note:  $\frac{\Gamma(\alpha-1)}{\Gamma(\alpha)} = \frac{(\alpha-2)!}{(\alpha-1)!} = \frac{1}{\alpha-1}$ . Since  $E(X) = \alpha\beta$ ,  $E\left(\frac{1}{X}\right)$  is not equal to  $\frac{1}{E(X)}$ . Thus,  $\frac{k^*}{T}$  is a biased estimator and  $E\left(\frac{k^*}{T}\right) = \lambda \left[ \frac{k^*}{k^*-1} \right]$ . The unbiased estimator is  $\frac{k^*-1}{T}$  and its variance is  $\text{Var}\left(\frac{k^*-1}{T}\right) = \frac{\lambda^2}{k^*-2}$ .

In Estimation Procedure 2, notice that  $t^*$  is fixed and  $K$  is a Poisson random variable. Since the likelihood function of  $K$  is again proportional to  $\lambda^k e^{-\lambda t^*}$ , the MLE is  $\frac{K}{t^*}$ . Now,  $E\left(\frac{K}{t^*}\right) = \lambda$ . Obviously, this estimator is unbiased and its variance is  $\text{Var}\left(\frac{K}{t^*}\right) = \frac{\lambda}{t^*}$ . If we set  $\text{Var}\left(\frac{K}{t^*}\right) = \text{Var}\left(\frac{k^*-1}{T}\right)$ , then  $\frac{\lambda}{t^*} = \frac{\lambda^2}{k^*-2}$  or  $\lambda = \frac{k^*-2}{t^*}$ . So if  $k^*$  and  $t^*$  are selected such that  $\lambda$  is approximately equal to  $\frac{k^*-2}{t^*}$ , then the two procedures are considered equivalent.

Which estimator is preferable? If  $k^*$  is large, then Estimation Procedure 1 is essentially unbiased. For a very small number of occurrences, however, this estimated demand rate may be too biased to be useful. Now for products with very low sales, selecting a large  $t^*$  may not be practical and selecting a large  $k^*$  may take much longer than is reasonable in observing the required number of units sold. Therefore, inventory managers should be aware of these issues when choosing estimators to forecast demands.

For products with high demand, Estimation Procedure 1 is preferable since the variance of the demand rate estimate quickly declines for large  $k^*$ . In this study, however, Estimation Procedure 2 will be used since otherwise an excessively large  $T$  may be required for a fixed  $k^*$ . Moreover, if a small  $k^*$  is

used, the variance for Estimator 1 could be large and the resulting demand rate estimate might be rather biased.

#### Estimating Demand Rate for a Family of Products with Zero Demand

Assume that the demand for each product in a large pool of  $n$  products follows a Poisson distribution. The underlying unknown demand rate of product  $i$  is  $\lambda_i$ ,  $i = 1, 2, \dots, n$ . When products sell at a reasonable rate, we can use either of the previously discussed procedures to estimate the future demand. However, for products showing no sales over a specified period, a predicted sales rate of zero for future sales may be too conservative as many products can still sell given a sufficiently long period of time. Consider Estimation Procedure 1. The value of  $k^*$  has to be at least 3 to estimate the variance, which might not be practical since the resulting large value of  $T$  may preclude a timely decision to liquidate the slow-moving products. As for Estimation Procedure 2, determining the value to set for  $t^*$  is not an easy task for slow-moving products, as an inappropriate choice of  $t^*$  may result in no sales observed. Consequently, reliable estimates of future demand can be difficult to derive and it is clear that new techniques are needed.

To motivate the approach to slow-moving products used in this dissertation, an analogy can be made between estimating the future occurrence rate of product sales after exhibiting no sales over a specified time period and estimating the future error occurrence rate in newly released software that has been tested and all known bugs have been removed. The error rate of the undiscovered software bugs is unknown just as the future demand rate of the non-selling products is unknown. Ross (1985a) also assigns a value to the

probability of detecting a bug. In this work, sales are easy to detect, so this probability is one, since all sales are typically recorded. Therefore, the focus of this section is to suggest an estimator of the future demand rate for products with no sales history by adapting a methodology used in software reliability research.

According to Ross (1985a, 1985b, 2002), the estimated error rate of a software package after debugging is equal to the number of bugs that caused exactly one error,  $M_1(t)$ , divided by the length of the testing period  $t$  in which these errors occur, that is,  $\frac{M_1(t)}{t}$ . An attractive feature of this estimator is that the total number of bugs in the software package does not have to be known. In addition, it is possible for the error rate of each bug to be different although knowledge of the number of errors caused by each bug is required. The random variable  $M_1(t)$  represents the number of products that have sold only one unit in the context of slow-moving inventory.

Ross (2002) also provides the expected value of the square of the difference between the population demand rate and  $\frac{M_1(t)}{t}$ , which is equivalent to the expected value of  $\frac{M_1(t) + 2M_2(t)}{t^2}$  with  $M_2(t)$  denoting the number of products that have exactly two units of demand. Using this expression as an estimator has the advantage of not requiring knowledge of the total number of products in a large offering of a variety of available products. However, a disadvantage is that the distributions of  $M_1(t)$  and  $M_2(t)$  are not known. A normal approximation may be used, but the accuracy of this procedure may be dependent on the length of the observed time period and the number of products in the inventory system.

To more formally discuss the proposed estimator of the future demand rate of a pool of products that have not sold, a mathematical development is now presented. The results here are more detailed than those reported in Chapter 2 and are based on the assumption of independent Poisson demand processes for the products. Define  $\Psi_i(t) = 1$  if product  $i$  with demand rate  $\lambda_i$  has not sold by time  $t$  and 0 otherwise,  $i = 1, 2, \dots, n$ . The objective is to develop prediction intervals for the value of the random variable  $\Lambda(t)$  in equation (2).

$$\Lambda(t) = \sum_{i=1}^n \lambda_i \Psi_i(t) \quad (2)$$

If an observable random variable can be found whose expected value is the same as that of  $\Lambda(t)$ , then it can be used as an estimator. Note that:

$$E[\Lambda(t)] = \sum_{i=1}^n \lambda_i E[\Psi_i(t)] = \sum_{i=1}^n \lambda_i e^{-\lambda_i t} \quad (3)$$

Interestingly,  $\frac{M_1(t)}{t}$  and  $\Lambda(t)$  have the same expected value according to Ross (2002).  $\frac{M_1(t)}{t}$  is an attractive estimator since it can be easily computed in practice. For it to be a good estimator of  $\Lambda(t)$ , its difference with  $\Lambda(t)$  should be small. The following results demonstrate that  $\frac{M_1(t)}{t}$  is an unbiased estimator of  $\Lambda(t)$  and that the variance of the difference, or equivalently the expected squared difference, between the estimator and the unknown population rate  $\Lambda(t)$  decreases over time.

$$\begin{aligned}
E[M_1(t)] &= \sum_{i=1}^n \lambda_i t e^{-\lambda_i t} \\
E\left[\frac{M_1(t)}{t}\right] &= E[\Lambda(t)] \\
E[M_2(t)] &= \frac{1}{2} \sum_{i=1}^n (\lambda_i t)^2 e^{-\lambda_i t} \\
E\left\{\left[\Lambda(t) - \frac{M_1(t)}{t}\right]^2\right\} &= \sum_{i=1}^n \left(\frac{\lambda_i^2 e^{-\lambda_i t} + \lambda_i e^{-\lambda_i t}}{t}\right) \\
&= \frac{E[M_1(t) + 2M_2(t)]}{t^2}
\end{aligned} \tag{4}$$

As an extension to the above approach, an estimator of the future demand rate of a pool of products experiencing no more than one unit sold is proposed.

Define  $\Delta(t) = \sum_{i=1}^n \lambda_i I_i(t)$  as the future unknown demand rate for products having

exactly one unit of sale by time  $t$  where  $I_i(t) = 1$  if product  $i$  with demand rate  $\lambda_i$  has exactly one unit of sale by time  $t$  and 0 otherwise,  $i = 1, 2, \dots, n$ . The future demand rate for products with sales of no more than one unit is the sum of the random variables  $\Lambda(t)$  and  $\Delta(t)$ . The proposed estimator of  $\Lambda(t) + \Delta(t)$  is

$\frac{M_1(t) + 2M_2(t)}{t}$ . This estimator is unbiased since its expected value is  $\sum_{i=1}^n (\lambda_i e^{-\lambda_i t} +$

$\lambda_i^2 t e^{-\lambda_i t})$ , the same as the expected value of  $\Lambda(t) + \Delta(t)$ . An unbiased estimator of

the squared difference between  $\Lambda(t) + \Delta(t)$  and  $\frac{M_1(t) + 2M_2(t)}{t}$  is

$\frac{M_1(t) + 2M_2(t) + 6M_3(t)}{t^2}$  since the expected value of either is  $\sum_{i=1}^n (t^{-1} \lambda_i e^{-\lambda_i t} + t \lambda_i^2 e^{-\lambda_i t} +$

$t^2 \lambda_i^3 e^{-\lambda_i t})$ . If this expected squared difference is small, then  $\frac{M_1(t) + 2M_2(t)}{t}$  is a

reasonable estimator of  $\Lambda(t) + \Delta(t)$ . The reliability of a proposed prediction interval for  $\Lambda(t) + \Delta(t)$  will be examined in Chapter 4.

Note that if the demand rate is large, then  $\lambda_i e^{-\lambda_i t}$  is small for  $i = 1, 2, \dots, n$ . In this case, the expected value of  $M_1(t)$  implies a small number of products with only one unit of sale. This is consistent with what one would expect  $\frac{M_1(t)}{t}$  to be when products are selling at a fast rate. In addition, note that if  $t$  is large,  $\frac{M_1(t)}{t}$  will be small, thus implying a low demand rate for products that have not experienced a sale by time  $t$ . Equation (10) shows  $\frac{M_1(t) + 2M_2(t)}{t^2}$  to be an unbiased estimator of the squared difference between  $\Lambda(t)$  and  $\frac{M_1(t)}{t}$ . Since the expected value of this estimator is  $\sum_{i=1}^n \left( \frac{\lambda_i^2 e^{-\lambda_i t} + \lambda_i e^{-\lambda_i t}}{t} \right)$ , the expected squared difference becomes small when  $t$  is large.

A natural question to ask is: How large should  $n$  and  $t$  be so that the squared difference between  $\Lambda(t)$  and  $\frac{M_1(t)}{t}$  is sufficiently small? To address this question, a bound on  $E\left[\left(\Lambda(t) - \frac{M_1(t)}{t}\right)^2\right]$  is derived as follows. First, the value of  $\lambda_i$  that maximizes  $E\left[\left(\Lambda(t) - \frac{M_1(t)}{t}\right)^2\right]$  is computed.

$$f(\lambda_i) = \left(\frac{\lambda_i^2 + \lambda_i}{t}\right) e^{-\lambda_i t}, \quad i = 1, 2, 3, \dots, n$$

$$\text{To find the max of } f(\lambda_i), \quad \frac{df(\lambda_i)}{d\lambda_i} = 0. \quad (5)$$

$$\left(\frac{2\lambda_i + 1}{t}\right) e^{-\lambda_i t} + \left(\frac{\lambda_i^2 + \lambda_i}{t}\right) (-t) e^{-\lambda_i t} = 0$$

Therefore,  $t\lambda_i^2 + (t-2)\lambda_i - 1 = 0$  and  $\lambda_i = \frac{2-t+\sqrt{t^2+4}}{2t}$ ,  $i = 1, 2, \dots, n$ . Now, the

second derivative is presented.

$$\frac{d^2f(\lambda_i)}{d\lambda_i^2} = \frac{2}{t} e^{-\lambda_i t} - t \frac{2\lambda_i + 1}{t} e^{-\lambda_i t} - (2\lambda_i + 1) e^{-\lambda_i t} + t(\lambda_i^2 + \lambda_i) e^{-\lambda_i t}$$

$$\frac{d^2f(\lambda_i)}{d\lambda_i^2} = e^{-\lambda_i t} \left[ \frac{2}{t} + t\lambda_i^2 + (t-4)\lambda_i - 2 \right] \quad (6)$$

The second derivative evaluated at  $\lambda_i = \frac{2-t+\sqrt{t^2+4}}{2t}$  is

$$\begin{aligned} &= e^{-\lambda_i t} \left[ \frac{2}{t} + t \left( \frac{2-t+\sqrt{t^2+4}}{2t} \right)^2 + (t-4) \left( \frac{2-t+\sqrt{t^2+4}}{2t} \right) - 2 \right] \\ &= e^{-\lambda_i t} \left[ \frac{2}{t} + \frac{4-4t+t^2+2(2-t)\sqrt{t^2+4}+t^2+4}{4t} + \frac{2t-t^2-8+4t+t\sqrt{t^2+4}-4\sqrt{t^2+4}}{2t} - 2 \right] \\ &= e^{-\lambda_i t} \left[ \frac{2}{t} + \frac{4-2t+t^2+2\sqrt{t^2+4}-t\sqrt{t^2+4}}{2t} + \frac{6t-t^2-8+t\sqrt{t^2+4}-4\sqrt{t^2+4}}{2t} - 2 \right] \\ &= e^{-\lambda_i t} \left[ -\frac{\sqrt{t^2+4}}{t} \right] \end{aligned}$$

The expected squared difference  $E \left[ \left( \Lambda(t) - \frac{M_1(t)}{t} \right)^2 \right]$  is maximized at

$$\lambda_i = \frac{2-t+\sqrt{t^2+4}}{2t} \quad i = 1, 2, \dots, n. \text{ For } t \geq 2, \text{ a bound for this expression is}$$

established in (7).

$$\begin{aligned}
& \mathbb{E} \left[ \left( \Lambda(t) - \frac{M_1(t)}{t} \right)^2 \right] \\
&= \sum_{i=1}^n \left( \frac{\lambda_i^2 + \lambda_i}{t} \right) e^{-\lambda_i t} \\
&\leq \sum_{i=1}^n \left( \frac{\left( \frac{2-t+\sqrt{t^2+4}}{2t} \right)^2 + \frac{2-t+\sqrt{t^2+4}}{2t}}{t} \right) e^{-\frac{2-t+\sqrt{t^2+4}}{2t}t} \\
&= \sum_{i=1}^n \left( \frac{\left( \frac{4-2t+t^2+4\sqrt{t^2+4}-2t\sqrt{t^2+4}+t^2+4+4t-2t^2+2t\sqrt{t^2+4}}{4t^2} \right)}{t} \right) e^{-\frac{2-t+\sqrt{t^2+4}}{2t}t} \\
&= \frac{n}{t^2} \frac{(t+2\sqrt{t^2+4}+4)}{2t} e^{-\frac{2-t+\sqrt{t^2+4}}{2}} \\
&\leq 1.07 \frac{n}{t^2} \text{ assuming } t \geq 2.
\end{aligned} \tag{7}$$

Inequality (7) follows by noting that  $e^{-\frac{2-t+\sqrt{t^2+4}}{2}}$  is an increasing function in  $t$  with an upper bound of  $e^{-1}$  and that  $\frac{t+2\sqrt{t^2+4}+4}{2t}$  is a decreasing function in  $t$ .

The value of 1.07 is equal to the product of  $e^{-1}$  and the expression  $\frac{t+2\sqrt{t^2+4}+4}{2t}$  evaluated at  $t=2$ . The value of  $t$  was assumed to be at least two so that the constant 1.07 was slightly smaller and close to a round number like 1.0. This bound provides the practitioner with an upper limit for the expected squared error with knowledge of only  $n$  (number of total products observed) and  $t$  (time frame for observing sales). Thus,  $n$  and  $t$  can be selected in practice to provide a bound on the expected squared error.

To investigate the empirical distribution of  $\frac{M_1(t)}{t}$ , a simulation of 300 products with a mean time between sales of 40 hours (i.e.,  $\lambda_i = \frac{1}{40}$ ,  $i = 1, 2, \dots, 300$ ) was replicated 500 times over time periods of 40, 70, 130, and 150 hours to illustrate the changing behavior of  $\frac{M_1(t)}{t}$  as the length of the time period increases for a slow-moving product.

For the time period with  $t = 40$ , the skewness is -0.30085 for the distribution shown in Figure 1. Observe that in Figure 2 with  $t = 70$ , the skewness is negative but small in magnitude. In Figure 3 with  $t = 130$ , the skewness is small in magnitude and positive. Finally, in Figure 4 with  $t = 150$ , the skewness is approximately 0.31. Thus, the skewness of the estimator  $\frac{M_1(t)}{t}$  is somewhat negative for shorter time periods and eventually becomes positive for longer time periods.

The graphs in Figures 1 through 4 show that prediction intervals developed for the future demand rate may not be reliable if the actual distribution of  $\frac{M_1(t)}{t}$  is too skewed. Proposed prediction intervals with  $100(1 - \alpha)\%$  confidence

will be of the form  $\frac{M_1(t)}{t} - Z_{\alpha/2} \sqrt{\frac{M_1(t) + 2M_2(t)}{t^2}}$  to  $\frac{M_1(t)}{t} + Z_{\alpha/2} \sqrt{\frac{M_1(t) + 2M_2(t)}{t^2}}$  for

products that show no sales. The reliability of such intervals will be assessed in detail over a variety of demand rates and numbers of products observed in Chapter 4.

Figure 1. Distribution with no sales over a period of 40 hours.




Figure 2. Distribution with no sales over a period of 70 hours.




Figure 3. Distribution with no sales over a period of 130 hours.





Figure 4. Distribution with no sales over a period of 150 hours.



## Bayesian Methodology for Inventory Modeling

The Bayesian approach is different from the traditional Poisson approach in that it uses a prior distribution for the demand rate  $\lambda$ . Popovic (1987) proposed a Bayesian model in which the demand has a Poisson distribution, but an a priori gamma distribution  $\Gamma(\alpha, \beta)$  with the probability density function in (8) is assigned/assumed since  $\lambda$  is unknown.

$$f(\lambda; \alpha, \beta) = \frac{\beta^\alpha \lambda^{\alpha-1} e^{-\beta\lambda}}{\Gamma(\alpha)}, \alpha > 0, \beta > 0, \lambda > 0 \quad (8)$$

Assuming  $X_t$  is the random variable representing demand during the time interval  $[0, t]$ , the unconditional distribution of demand is displayed in (9).

$$P\{X_t = k\} = \int_0^\infty \frac{(\lambda t)^k}{k!} e^{-\lambda t} f(\lambda) d\lambda = \binom{\alpha + k - 1}{k} \left[ \frac{\beta}{\beta + t} \right]^\alpha \left[ \frac{t}{(\beta + t)} \right]^k, \alpha > 0, \beta > 0, t > 0 \quad (9)$$

where  $k = 0, 1, \dots$  and the notation  $\binom{n}{x}$  is the number of combinations of taking  $x$

items from  $n$  distinct items at a time. Note that the combination  $\binom{\alpha + k - 1}{k}$  in (9)

can also be written equivalently as  $\binom{\alpha + k - 1}{\alpha - 1}$ . Thus,  $X_t$  has a negative binomial

distribution  $NB(\alpha, \frac{\beta}{\beta + t})$ . Popovic (1987) shows that denoting  $P\{X_t = k\}$  by  $p_k$ ,

where  $k = 0, 1, \dots$ , we have

$$p_0 = \left[ \frac{\beta}{\beta + 1} \right]^\alpha, \quad (10)$$

$$p_k = \left[ \frac{\alpha + k - 1}{k(\beta + 1)} \right] p_{k-1}$$

By applying Bayes' rule for the first time period, we have:

$$f(\lambda | X_1) = \frac{P(X_1 | \lambda)f(\lambda)}{\int_0^{\infty} P(X_1 | \lambda)f(\lambda)d\lambda} \quad (11)$$

The a posteriori distribution of  $\lambda$  over the first unit time interval  $I_1 = [0, 1]$  is:

$$f(\lambda | X_1) = \frac{(\beta + 1)^{\alpha + X_1} \lambda^{(\alpha + X_1 - 1)} e^{-(\beta + 1)\lambda}}{\Gamma(\alpha + X_1)} \quad (12)$$

The a posteriori distribution is also illustrated for the next time to show a pattern in the form of the distribution and is displayed in (13) for time interval  $I_2 = [1, 2]$ .

$$\begin{aligned} P\{X_2 = k\} &= \int_0^{\infty} \frac{\lambda^k e^{-\lambda}}{k!} f(\lambda | X_1) d\lambda \\ &= \left[ \frac{\Gamma(\alpha + k - 1)}{k!} \Gamma(\alpha + X_1) \right] \left[ \frac{\beta + 1}{\beta + 2} \right]^{\alpha + X_1} \left[ \frac{1}{\beta + 2} \right]^k \end{aligned} \quad (13)$$

Thus  $X_2 \sim \text{NB}(\alpha + X_1, \frac{\beta + 1}{\beta + 2})$ . It then follows that the a posteriori distribution of  $\lambda$

for the second time interval  $I_2$  after demand  $X_2$  occurs will be:

$$\lambda | X_1, X_2 \sim \Gamma(\alpha + X_1 + X_2, \beta + 2)$$

Furthermore, it can be shown that the general a posteriori distribution of  $\lambda$  is

$$\lambda | X_1, X_2, \dots, X_n \sim \Gamma(\alpha + \sum_{i=1}^n X_i, \beta + n) \quad (14)$$

After observing  $X_1, X_2, \dots, X_n$ , one can prove that (15) is the distribution of demand at interval  $I_{n+1}$ .

$$X_{n+1} \sim \text{NB}\left(\alpha + \sum_{i=1}^n X_i, \frac{\beta + n}{\beta + n + 1}\right) \quad (15)$$

Popovic (1987) states that the optimal inventory level in the  $i^{\text{th}}$  unit time interval  $I_i$  can be computed by considering the ratio of the cost of surplus  $C_1$  and

the cost of shortage  $C_2$  of an item as well as the demand described by the a posteriori distribution. As additional information is accumulated, the a posteriori distribution of demand can be updated to improve the accuracy of the parameter estimates. The optimal inventory level  $r_i^*$  for product  $i$  will be such that it satisfies one of the inequalities in (16).

Bayes Model Approach:

$$\sum_{k=0}^{r_i^*-1} \binom{\alpha+k-1}{k} \left(\frac{\beta}{\beta+1}\right)^\alpha \left(\frac{1}{\beta+1}\right)^k < \frac{C_2}{C_1+C_2} \leq \sum_{k=0}^{r_i^*} \binom{\alpha+k-1}{k} \left(\frac{\beta}{\beta+1}\right)^\alpha \left(\frac{1}{\beta+1}\right)^k$$

Poisson Model Approach:

(16)

$$\sum_{k=0}^{r_i^*-1} e^{-\lambda t} \frac{(\lambda t)^k}{k!} < \frac{C_2}{C_1+C_2} \leq \sum_{k=0}^{r_i^*} e^{-\lambda t} \frac{(\lambda t)^k}{k!}$$

Applications that use the ratio of  $C_1$  and  $C_2$  to determine optimal inventory levels include style goods and perishable items (Silver et al., 1998). Finding appropriate costs for shortage and surplus goods for inventory with intermittent demand may not always be practical and, thus, could potentially be a limitation to this approach. Hill (1999) states that Popovic's (1987) approach is a Bayesian treatment of an essentially single-period model. Popovic (1987) repeats his approach over several periods. Popovic (1987) describes an application of this approach to the ordering of spare parts from a warehouse at the beginning of each month and explains how the equalities in (16) can be used for ordering optimal stock levels.

The research in this dissertation examines the performance of Popovic's (1987) approach under various demand rates, but does not investigate alternative model formulations that may be more applicable to managing

inventory items such as spare parts over multi-periods. Several papers in the literature remark that single-period based inventory models can be extended. Matsuyama (2006) generalizes the so-called newsboy model and allows products to be unsold at the end of a period. Koulamas (2006) examines the newsvendor problem and remarks that the single-period assumption can be lifted as long as the multi-period demands are independent, identically distributed random variables. Their paper suggests the use of present values of cash flows in their model to allow for the extension to multi-periods.

DuMouchel (1999) extends the Bayes procedure introduced by Popovic (1987) to a two-group approach, but it is not examined here. The basic change is to assume that the demand rate can come from one of two a priori probability distributions. In other words, the a priori probability density of  $\lambda$  can be written as  $f(\lambda; \alpha_1, \beta_1, \alpha_2, \beta_2, p) = p\Gamma(\lambda; \alpha_1, \beta_1) + (1 - p)\Gamma(\lambda; \alpha_1, \beta_1, \alpha_2, \beta_2)$ ,  $\alpha_1 > 0$ ,  $\beta_1 > 0$ ,  $\alpha_2 > 0$ ,  $\beta_2 > 0$ ,  $0 \leq p \leq 1$ .

### Stopping Rules

A stopping rule that is a function of future demand rate can be applied in deciding to continue or discontinue selling a product. For slow-moving merchandise, an effective rule may be difficult to establish. Retail managers may re-evaluate the future demand rate of a group of non-selling or slow-moving products at the end of specified time periods. An easy-to-implement stopping rule can be used to discontinue products whose projected demand rate is less than a threshold value. Typically, such a threshold is based on financial considerations.

If the upper endpoint of a one-sided prediction interval for the product's future demand rate is below the threshold, for instance, then the decision is to liquidate.

This study will assess the robustness of the reliability of a one-sided prediction interval for the future demand rate across a variety of parameters. The one-sided prediction interval is the same as the proposed two-sided interval discussed previously except the  $\alpha$  level (Type I error) is not divided by 2 and the lower endpoint is not estimated. The proposed one-sided prediction interval is

$$\left[0, \frac{M_1(t)}{t} + Z_\alpha \sqrt{\frac{M_1(t) + 2M_2(t)}{t^2}} \right].$$

Stopping rules that are applicable to debugging in software development have been proposed by Forman and Singpurwalla (1977) and Ross (1985b).

Ross (1985b) proposes a stopping rule for software development based on the estimated error rate. This same procedure can be adapted and applied to slow-moving inventory. One possible procedure is to use a quality control-type rule that uses an upper bound similar to that of a confidence interval. Assume that  $\varepsilon(t)$  is the estimate of  $\Lambda(t)$ , which is the true demand rate. The estimator

$\frac{M_1(t)}{t}$  could be used as a substitute for  $\varepsilon(t)$ . The following guideline is adapted

from Ross (1985b) for discontinuing holding certain inventory:

$$\varepsilon(t) + 3\sqrt{E[\varepsilon(t) - \Lambda(t)]^2} < A, \text{ where } A \text{ is the minimal acceptable demand rate}$$

determined by management. The reliability of this procedure will not be examined in this study. However, the performance of the one-sided prediction intervals will provide insights into the feasibility of this guideline.

## Research Questions

Research Question 1: For intermittent-demand products, can reliable two-sided prediction intervals be derived using the  $\frac{M_1(t)}{t}$  estimator (Ross, 2002) to forecast the demand rate of products that have not sold over a specified time period?

This research question is addressed via a Monte Carlo simulation study. Empirical confidence levels are evaluated against their nominal  $100(1 - \alpha)\%$  levels across a variety of conditions of  $n$  (number of products),  $\lambda$  (demand rates – both homogeneous and mixed), and  $t$  (length of time period). For each of the research questions, the demand for one product is assumed to be independent of the demand for another product as this is a standard assumption in many applications in the literature. Thus, demand will be simulated independently for each product. In addition, this study does not distinguish between sales and demand. Thus, all demand is assumed to result in a sale. In practice, this assumption does not necessarily hold.

Research Question 2: For intermittent-demand products, can reliable one-sided prediction intervals be derived using the  $\frac{M_1(t)}{t}$  estimator (Ross, 2002) to forecast the demand rate of products that have not sold over a specified time period to be incorporated in a stopping rule procedure?

This research question will be addressed in a similar fashion to that of Research Question 1. Through a Monte Carlo simulation study, the reliability of upper one-sided prediction intervals is assessed.

Research Question 3: For intermittent-demand products, can reliable prediction intervals be developed using an extension of the estimators assessed in Research Questions 1 and 2 to forecast the demand rate of products that have sold no more than one unit over a specified time period?

The prediction intervals proposed for Research Questions 1 and 2 are extended to the case in which products have sold no more than one unit over a specified time period. The extensions of these estimators were discussed earlier. As for the previous research questions, a Monte Carlo simulation is conducted to evaluate the reliability of the proposed prediction intervals.

Research Question 4: How effective is the Bayesian approach to estimating optimal inventory levels for moderate-demand products as compared to one using a maximum likelihood estimator of the demand rate parameter of a Poisson process?

As described previously in this chapter, Popovic (1987) provides an estimation procedure for the optimal number of inventory items to stock using a Bayesian approach. The costs of surplus and shortage determine such an inventory level. This approach is compared with a traditional approach in which a Poisson distribution is assumed for the demand rate estimated from the previous period's data. A Monte Carlo simulation analysis of these estimators is performed to determine the effectiveness of these procedures over sequential time intervals.

Research Question 5: How effective is the Bayesian approach to estimating optimal inventory levels for intermittent-demand products as compared to one

using a maximum likelihood estimator of the demand rate parameter of a Poisson process?

This question addresses the performance of the Bayesian approach to estimating optimal inventory levels just as described in Research Question 4 but considers products having a lower demand rate (intermittent demand).

Research Question 6: How effective is the Bayesian approach using a mixture of prior distributions to estimate optimal inventory levels for intermittent-demand products as compared to one using a maximum likelihood estimator of the demand rate parameter of a Poisson process?

In Research Question 5, the prior distribution for the demand rate is assumed to be a gamma distribution. In addressing this question, however, the prior distribution of the demand rate will be the same as that used in the previous question or the estimate is the predicted demand rate used in Research Question 3, if the product experiences one unit of sales or no sales at all. A simulation-based comparison of three estimators (i.e., traditional method with Poisson distribution and estimated rate from observed data, Bayesian approach as described in Popovic (1987), and Bayesian approach with mixture of two prior groups) is made over a variety of parameter values for the demand rate.

Research Question 7: Are the proposed prediction intervals investigated in Research Questions 1, 2, and 3 reliable for predicting future demand rates of slow-moving products using data from a major national retailer?

The reliability of the prediction intervals used in addressing Research Questions 1, 2, and 3 will be examined using real-world sales data from a large national retailer. More specifically, weekly sales data for 103 weeks over a 2-year period for approximately 600 slow-moving SKUs will be used in assessing the performance of the proposed prediction intervals.

## CHAPTER 4

### ANALYSIS AND RESULTS

This chapter presents results of the analysis of the seven research questions posed in Chapter 3. This chapter is organized by research question. Each research question is restated, the procedure for analyzing the question is discussed, and results are presented, mostly by displaying graphs of the performance of proposed methods.

The analysis of the performance of models proposed in this dissertation will be conducted using a Monte Carlo simulation study. A researcher may ask, Why use simulation instead of an analytical approach? An analytical approach may be too involved considering the number of parameters that vary. The theoretical distributions of the proposed estimators depend on various parameters: time, demand rate, and number of products. The distributions of the estimators may not be of a form that is easy to implement. The proposed prediction intervals make use of a normal distribution approximation. To determine the robustness of these prediction intervals, a simulation study can assess the models across a variety of conditions, including mixtures of demand rates across product subgroups. Oftentimes, a simulation study is conducted and theoretical distributions are later derived to explain the behavior of proposed methods. Examples of studies that use Monte Carlo simulations to assess the Type I error rate of a confidence interval or hypothesis test are Zwick (1986) and Harwell (1991).

## Research Question 1

For intermittent-demand products, can reliable two-sided prediction intervals be derived using the  $\frac{M_1(t)}{t}$  estimator (Ross, 2002) to forecast the demand rate of products that have not sold over a specified sales period?

The performance of two-sided prediction intervals (TSPIs) for the future demand rate of slow-moving products with zero observed sales was assessed using a normal distribution approximation for the proposed  $\frac{M_1(t)}{t}$  estimator. A Monte Carlo simulation utilizing 5,000 replications was conducted to estimate the Type I error rate of the prediction intervals across experimental parameters, namely, demand rate and number of products, over a specified time frame. Thus, empirical confidence levels for TSPIs can be evaluated against their nominal  $100(1 - \alpha)\%$  confidence levels.

The term product group will be used to denote a group of different products, which may be dissimilar in features, but similar in demand rates. The term time frame will refer to the length of time in which demand for products is observed before computing the proposed prediction intervals. An implicit assumption of the simulation is that the demand for a product is independent of the demand for other products. Prediction intervals will be referred to as being robust if they maintain their nominal Type I error rate over a variety of experimental conditions.

An appropriate time frame needed to be selected to observe simulated demand. After preliminary experimentation, a time frame of 100 time units was

selected. This choice provided sufficient time to observe sales of slow-moving products over a variety of demand rates and still be able to assess the performance of the proposed prediction intervals. The time interval is also of interest to the national retailer who supplied data for this research. Furthermore, it should be noted that time units can be hours, days, weeks, months, or quarters. For example, the time frame of 100 units can approximate one calendar quarter (90 days) if the time unit is days or two years (104 weeks) if the time unit is weeks. The prediction intervals are used to forecast only the demand rate of products having zero observed demand, not the demand rate of the entire group of products. Sales and demand are assumed to be the same in this study. The demand rate will be expressed in terms of mean time between demands (MTBD).

#### Effect of Product Group Size on Prediction Intervals

An initial simulation was conducted to assess the effect of product group size on the reliability of the proposed prediction intervals. Product group sizes between 50 and 1,000 were selected and the effect of the group size on the performance of the prediction intervals was assessed at four demand levels in terms of MTBD: 100, 200, or 300, or mixture of 50 and 400. Product group sizes were increased by 50 up to 500 products and then by 100 up to 1,000 products. Product group sizes were incremented in this fashion to keep the total number of simulations reasonable while still gaining insight into the performance of the models when product size is large. Product group sizes under 500 were considered more practical to investigate than group sizes over 500. Throughout the simulation, product group sizes and MTBD will not be evenly spaced so that

simulations can reveal information over wider ranges of parameters while still examining performance over parameter values where performance is thought to be changing quickly. The product mixture of 50 and 400 MTBDs consisted of 25 products having an MTBD of 400 and the remaining 25 to 975 products having an MTBD of 50.

Ideally, the proposed prediction intervals should have empirical Type I error rates near their nominal alpha values, i.e., near the alpha levels of 10%, 5%, and 1% for 90%, 95%, and 99% prediction intervals, respectively. An empirical Type I error will be considered near its nominal alpha value if this error is within plus or minus two standard deviations of the nominal value (Zwick, 1986; Harwell, 1991). For example, at the 99% confidence level, the empirical

Type I error for 5,000 simulations must be between  $0.01 \pm 2\sqrt{\frac{(0.01)(1-0.01)}{5,000}}$  or

from 0.007 to 0.013 for the prediction interval to be considered reliable. For the 95% and 90% confidence levels, these intervals are 0.044 to 0.056 and 0.092 to 0.108, respectively. These end values will be referred to as the reliability

boundaries. Since the proposed prediction intervals rely on a normal approximation, an appropriate product group size must be determined for prediction intervals to be reliable. Figure 5, graphically, and Table 7, numerically, reveal that the empirical Type I error rates of the proposed TSPIs over four demand levels and 15 product group sizes are mostly within the reliability boundaries for group sizes over 200. Actual values for the empirical Type I errors will generally not accompany the graphs in the presentation of the simulations since trends and behavior patterns can be deduced from the plots.

Figure 5. Empirical Type I error versus product group size for TSPI.




Table 7

Empirical Type I Error for TSPI in Figure 5

Number of Products	90% TSPI with MTBD = 100	90% TSPI with MTBD = 200	90% TSPI with MTBD = 300	90% TSPI with MTBDs = 400/50
50	0.103	0.110	0.118	0.112
100	0.107	0.105	0.117	0.105
150	0.103	0.104	0.104	0.108
200	0.099	0.102	0.101	0.104
250	0.106	0.102	0.104	0.103
300	0.104	0.107	0.107	0.096
350	0.102	0.104	0.106	0.104
400	0.094	0.100	0.104	0.100
450	0.100	0.098	0.110	0.100
500	0.101	0.109	0.111	0.106
600	0.100	0.100	0.109	0.102
700	0.103	0.094	0.102	0.102
800	0.103	0.095	0.100	0.100
900	0.106	0.099	0.106	0.100
1,000	0.104	0.094	0.104	0.102
Number of Products	95% TSPI	95% TSPI	95% TSPI	95% TSPI
50	0.055	0.061	0.069	0.058
100	0.052	0.054	0.060	0.053
150	0.047	0.054	0.054	0.053
200	0.046	0.054	0.051	0.051
250	0.048	0.055	0.056	0.054
300	0.051	0.053	0.058	0.048
350	0.049	0.051	0.054	0.053

Table 7 Continued

Empirical Type I Error for TSPI in Figure 5

Number of Products	95% TSPI	95% TSPI	95% TSPI	95% TSPI
400	0.045	0.050	0.054	0.049
450	0.048	0.051	0.057	0.049
500	0.048	0.054	0.056	0.052
600	0.049	0.051	0.054	0.053
700	0.048	0.043	0.049	0.049
800	0.048	0.047	0.051	0.049
900	0.053	0.046	0.054	0.050
1,000	0.050	0.048	0.052	0.050
Number of Products	99% TSPI	99% TSPI	99% TSPI	99% TSPI
50	0.012	0.020	0.022	0.019
100	0.010	0.015	0.016	0.012
150	0.010	0.016	0.014	0.012
200	0.009	0.011	0.013	0.010
250	0.011	0.011	0.012	0.010
300	0.009	0.011	0.012	0.013
350	0.009	0.011	0.012	0.011
400	0.009	0.012	0.010	0.010
450	0.008	0.011	0.013	0.008
500	0.007	0.008	0.011	0.009
600	0.009	0.010	0.010	0.011
700	0.010	0.011	0.010	0.009
800	0.010	0.009	0.010	0.009
900	0.007	0.010	0.010	0.010
1,000	0.009	0.010	0.012	0.008

In Figure 5, for product group sizes that are small (i.e., 50 or 100) the empirical Type I error is higher than the nominal. Whether demand rates are either MTBD = 100, MTBD = 200, or mixed, the prediction intervals maintain nominal alpha levels for products group sizes of 200 or larger. That is, at these product group sizes, the empirical Type I error is within the reliability boundaries.


Overall, there are more empirical Type I error rates beyond the reliability boundaries at an MTBD of 300 than at the other three selected MTBDs. For an MTBD of 300, prediction intervals are not reliable for small product group sizes below 200 as well as for product group sizes of 450 and 500 at the 90% confidence level and for the product group size of 300 at the 95% confidence level. Furthermore, at an MTBD of 100, the empirical Type I errors of the TSPIs are close to their nominal values and are never outside of the reliability boundaries for the conditions selected in Figure 5. In Figure 5, the empirical Type I error rates generally trend toward the nominal alpha levels as the product group size increases, suggesting that for product group sizes that are large prediction intervals are more robust. However, very large product groups may not always be available in practice. A product group size of 200 appears to be the smallest size for a group of products that generally yields robust prediction intervals. In addition, a group of 200 products is feasible for many retailers to monitor, particularly for the national retailer providing data for this study. Thus, a size of 200 is selected as a reference point for the size of product groups analyzed here. Under certain conditions, such as for a mixture of MTBDs, a smaller product group size may yield robust prediction intervals for the demand rates considered.

## Effect of Demand Rates on Two-Sided Prediction Intervals

The effects of demand rates on the reliability of the proposed TSPIs were examined. Since a product group size of 200 is considered a reference point as mentioned previously, a group of 200 products was used in further analysis in which MTBDs ranging from 15 to 1,000, as shown on the horizontal axis of Figure 6, were used to examine the reliability of the TSPIs. Figure 6 reveals that the prediction intervals' empirical Type I errors trend upward when the demand rate is relatively high (relatively short MTBD) or the demand rate is relatively low (relatively long MTBD). The 90% and 95% TSPIs have confidence levels close to their respective 90% and 95% nominal levels for MTBDs as high as 800. The 99% TSPIs are reliable for MTBDs between 30 and 300. The simulations suggest that the empirical Type I error rates trend closer to the nominal Type I error as the underlying MTBD increases from 15 to around 100, but then gradually trend upward as the MTBD increases toward 1,000 at which point the TSPIs are no longer reliable.

Not every possible interval in the range investigated was examined. For faster moving products, changes in the MTBD resulted in more substantial changes in the empirical Type I error rates of the TSPIs than for larger MTBDs so additional simulations were conducted in that range. Again, to keep the total number of simulations at a reasonable level, values for MTBD were selected so that more simulations were conducted at MTBDs that were not too extreme. A possible explanation for the trends illustrated in Figure 6 is that the TSPIs have

Figure 6. Empirical Type I error versus MTBD for TSPI for 200 products.



empirical Type I errors closer to their nominal value when there are a sufficient number of products with no sales and a sufficient number of products with some sales. When demand is high (short MTBD), few periods with no sales exist. As the demand decreases too much, many periods of zero demand exist, resulting in few products with observed sales. Interestingly, the 99% TSPI is reliable over a smaller range of MTBDs than the 90% and 95% TSPIs.

#### Two-Sided Prediction Intervals Using a Mixture of Demand Rates

Obviously, in practice, not all products have the same demand rate. A plausible scenario is that products can be categorized as either relatively faster moving or as relatively slower moving by managers. Figures 7 through 10 demonstrate the reliability of the proposed prediction intervals across different mixtures of two demand rates. A group of 200 products is divided into two subgroups, each with a different MTBD. In addition to varying the MTBD, we also change the subgroup sizes gradually increasing one from a minority group to a majority group. The purpose of this experimental design is to find out if the ratio of subgroup sizes – one subgroup consisting of relatively fast-moving products and one of relatively slow-moving products – affects the reliability of the TSPIs.

In Figure 7, the following pairs of MTBDs were assigned to the two subgroups: 100 and 10, 200 and 10, 400 and 10, or 1,000 and 10. Thus for this set of mixtures, the second subgroup of products has a relatively higher demand with an MTBD of 10. While an MTBD of 10 may be considered slow by some standards, for the purposes of this study, this rate is considered high. In Figure 7, the left side of the graph shows the empirical Type I error rate of TSPIs when the

vast majority of products (175 out of 200 products for the leftmost combinations of subgroup sizes) have an MTBD of 10. On the right of the graph, the empirical Type I error rate of TSPIs are shown when the majority of products have an MTBD of 100, 200, 400 or 1,000 and the rest have an MTBD of 10.

Figure 7 reveals that for a mixture of 1,000/10 MTBDs, the TSPIs are virtually not reliable. The mixture of 400/10 MTBDs is not much better. The empirical Type I error rates generally tend to be lower at the extreme right of the graph and higher on the left side. The proposed prediction intervals are not reliable if the number of relatively faster moving products is too high as compared to the number of slower moving products. In other words, the prediction intervals generally have empirical Type I errors closer to their nominal values when the number of products with an MTBD of 10 is minimized. For the mixture of 100/10 MTBDs, the TSPIs are reliable but tend to have higher empirical Type I error rates for subgroup sizes of 25 with an MTBD of 100 and 175 with an MTBD of 10. The TSPIs for the mixture of 200/10 MTBDs are reliable at the 90% confidence level except for the subgroup sizes 25 with an MTBD of 200 and 175 with an MTBD of 10. At the 95% and 99% confidence levels, these TSPIs are generally not reliable.

In Figure 8, subgroup sizes for the relatively slower moving products and relatively faster moving products were selected at the following MTBDs: 10/50, 100/50, 200/50, 400/50, or 1,000/50. Thus, the second subgroup of products has the relatively shorter MTBD of 50 except for the first mixture of 10/50.

Figure 7. Empirical Type I error for TSPI for mixtures of MTBDs with MTBD of 10 for second subgroup.




Figure 8. Empirical Type I error for TSPI for mixture of MTBDs with MTBD of 50 for second subgroup.




Figure 9. Empirical Type I error for TSPI for mixture of MTBDs with MTBD of 100 for second subgroup.




Figure 10. Empirical Type I error for TSPI for mixture of MTBDs with MTBD of 400 for second subgroup.




Figure 8 reveals that the reliability of the TSPIs is much improved from that in Figure 7. All empirical Type I error rates are within their reliability bounds except for the 10/50 MTBDs at the 99% confidence level and the 400/50 MTBDs at the 90% confidence interval. All 95% TSPIs in this figure are reliable.

In Figure 9, the following MTBD combinations are assigned: 10/100, 200/100, 400/100, or 1,000/100. The TSPIs are reliable for all combinations of subgroup sizes except the 175/25 combination. In this case, the proportion of products with an MTBD of 100 is the smallest. At the 175/25 subgroup sizes, the TSPIs with MTBDs = 10/100 mixture are not reliable for all three confidence levels. Figure 10 is constructed in a similar fashion to the previous graphs, except that the product group size of the second subgroup is 400. In it, the TSPIs for subgroups with very low or very high demand, i.e., with MTBDs = 10/400 or 1,000/40, are not reliable at the 95% and 99% confidence levels. The empirical Type I error rates for the MTBD mixture of 200/400 are all above their nominal Type I error rates, but are within their reliability boundaries with the exception of the 25/175 combination of subgroup sizes.

Figures 7 through 10 consistently illustrate that groups of products with an MTBD of 100 as one of the MTBDs for the two subgroups display yield TSPIs with their empirical Type I error rates closer to their nominal Type I error rates. Too many relatively high demand (MTBD = 10) products in combination with the slower selling products results in more TSPIs that are not reliable. Too many products with very slow demand rates (MTBD = 1,000) likewise tends to cause the proposed prediction intervals to not be robust with respect to maintaining

their nominal Type I error.

An important observation from viewing Figures 7 through 10 is that the TSPIs performing more consistently with respect to maintaining their nominal Type I error have an MTBD not far from 100. The results of these figures are consistent with the results in Figure 7. In that figure, small MTBDs and large MTBDs (outside of the MTBD range of 30 to 300) generally resulted in higher Type I error rates.

### Research Question 2

For intermittent-demand products, can reliable one-sided prediction intervals be derived using the  $\frac{M_1(t)}{t}$  estimator (Ross, 2002) to forecast the demand rate of products that have not sold over a specified time period to be incorporated in a stopping rule procedure?

A stopping rule procedure is any decision-making process in which a decision is required on whether to stop carrying a product or group of products. A stopping rule procedure can incorporate an estimate of the future demand rate of products into its analysis to determine the continuation of inventory. Upper-sided prediction intervals for forecasting future demand rate for products showing no sales over a specified time frame will be assessed for robustness with respect to their nominal Type I error rate. These intervals are constructed similar to the TSPIs used to answer Research Question 1 except that only the upper prediction interval endpoint is used with the corresponding confidence level.

Decision rules can be formed in which a group of products will be discontinued if there is a high degree of confidence that the future demand rate

will be below a certain threshold. The threshold is typically based on economic considerations, which will not be explored in this study. For example, see Horodowich (1979) for a model that considers cost of capital, income taxes, selling price inflation, scrap value, and magnitude of inventory on hand, when considering when to discontinue a product. The results of the analysis of this question will be used to assess the feasibility of using the upper endpoint of a one-sided prediction interval (OSPI) in a stopping rule. As presented in Chapter 3, the OSPI is the interval  $[0, \frac{M_1(t)}{t} + Z_\alpha \sqrt{\frac{M_1(t) + 2M_2(t)}{t^2}}]$ . Inventory managers may decide to stop holding a group of products if the estimated future sales rate, i.e., the upper endpoint of the OPSI, is below a threshold value.

As in the simulation study addressing Research Question 1, a period of 100 time units was selected as the time frame to collect data on sales of slow-moving products. A Monte Carlo simulation with 5,000 replications of the demand for a group of products over 100 units of time are performed to determine if OSPIs are reliable as a measure to compare with a threshold value for making critical decisions about a subgroup of non-selling products. Therefore, empirical Type I errors of OSPIs are assessed across a variety of conditions for the number of products and the demand rates.

In Figure 11, the number of products in a group ranges from 50 to 1,000 products across four demand rate levels for three confidence levels, which is the same as the parameters selected previously. In Figure 11, as the product group size increases the empirical Type I error of the OSPIs trends toward the nominal Type I error rates. This trend is consistent with the Two Sided Prediction

Intervals. A larger number of products appear to be necessary for most of the OSPIs to maintain their nominal Type I error rate. The OSPIs appear to be less stable when the demand rate is relatively low (i.e., MTBD = 300).

To further compare the behavior of the OSPIs with TSPIs, empirical Type I errors for the OSPIs are computed using a product group size 200 similar to that shown previously in Figure 6. The results are presented in Figure 12. The demand rate in terms of MTBD has a greater effect on the OSPIs than on the TSPIs at similar demand levels. For a product group size of 200 or smaller, the prediction intervals only perform marginally well with demand rates around MTBD = 100 and progressively worse when the demand increases above MTBD = 100 or decreases below MTBD = 100. Note that the number 100 for MTBD also coincides with the time frame used to observe sales of products.

Figure 11 and Figure 12 provide insight into the behavior of the OSPIs and allow for comparison with the TSPIs previously discussed in Figures 5 and 6. The figures suggest that the OSPIs are not as robust with respect to maintaining their nominal Type I error as the TSPIs. Generally, the OSPIs are reliable for higher demand rates (shorter MTBD) with relatively larger product group sizes. The OSPIs should be used with caution for product group sizes below 300 or with very low demand rates. Reliable OSPIs can be obtained for use in a stopping rule. The conditions under which OSPIs should be considered reliable include product group sizes that are large and an observed time period that should approximate or be close to the MTBD of the product group.

Figure 11 . Empirical Type I error versus product group size for OSPI.





Figure 12. Empirical Type I error versus MTBD for OSPI for 200 products.



### Research Question 3

For intermittent-demand products, can reliable prediction intervals be developed using an extension of estimators assessed in Research Questions 1 and 2 to forecast the demand rate of products that have sold no more than one unit over a specified time period?

Sales managers may need a forecast for not only the products that have not sold, but also products that had few sales. That is, a prediction interval for slow-moving products with less than some minimum number of observed sales over a specified time period may be desired. This prediction interval could be used to assess the feasibility of continuing to carry the slow-moving products. The prediction interval formulas used in addressing Research Questions 1 and 2 are extended to the case in which the future demand rate is estimated for products having no more than one sale. The procedure used can be extended again to estimate future demand rate for products with no more than two sales over a specified time frame. An extension would be to construct prediction intervals for slow-moving products having no more than some fixed number of sales.

As discussed in Chapter 3, the proposed estimator for the future sales rate of products having no more than one sale over a specified time frame is

$\frac{M_1(t) + 2M_2(t)}{t}$  since it is an unbiased estimator of the underlying demand rate.

Also discussed in Chapter 3 is the unbiased estimator for the expected squared difference of this estimator and the future demand rate. This estimator is

$\frac{M_1(t) + 2M_2(t) + 6M_3(t)}{t^2}$ . The end points of the proposed prediction interval for the

future demand rate of products having no more than one sale by time period  $t$  are

$\frac{M_1(t) + 2M_2(t)}{t} \pm Z_{\alpha/2} \sqrt{\frac{M_1(t) + 2M_2(t) + 6M_3(t)}{t^2}}$ . This prediction interval is an

extension of the prediction intervals proposed in the first two research questions.

The prediction intervals in Research Questions 1 and 2 will be referred to as the Zero Sales prediction intervals. That is, the Zero Sales prediction intervals

determine future demand rate for products exhibiting no sales over a specified

time frame. The proposed prediction interval addressing Research Question 3

will be referred to as the Zero and One Sales prediction interval. These prediction

intervals determine the future demand rate for products having no more than one

sale over a specified time frame.

A Monte Carlo simulation with 5,000 replications is conducted similar to that in examining Research Questions 1 and 2 to assess the proposed prediction interval's empirical Type I error rate. Since the Zero and One Sales prediction interval is estimating a demand rate for potentially more products than the Zero Sales prediction interval, this new prediction interval is expected to perform better with products having higher demand, since more demand history will be available to use in the model. To assess this characteristic, an initial comparison is made between these two prediction intervals. Then, the effect of the product group size is examined and finally the effect of the demand rate in terms of MTBD on the newly proposed prediction interval is studied. Similar to previous questions, only select MTBDs are reported.

Figure 13. Empirical Type I error for zero and one sales and zero sales OSPI.





Figure 14. Empirical Type I error for zero and one sales and zero sales TSPI.



The empirical Type I errors for the Zero Sales and for the Zero and One Sales prediction intervals over MTBD values ranging from 15 to 1,000 are presented in Figure 15 (OSPIs) and Figure 16 (TSPIs) for a group of 200 products. These results indeed reveal that for high demand rates, such as MTBD = 15 and MTBD = 20, that the Zero and One Sales prediction intervals have empirical Type I error rates closer to their nominal Type I error rates. This occurs because there is more information available at relatively higher demand rates for the Zero and One Sales prediction intervals. However, as the MTBD increases, the Zero and One Sales prediction interval does appear to have higher empirical Type I error rates than the Zero Sales prediction interval. The TSPIs are more robust in maintaining their nominal Type I error rate than the OSPIs. As noted previously, for MTBDs close to 100, the OSPIs and TSPIs are generally reliable.

Comparable to the graphs presented in addressing Research Question 1 and 2, the effect of changing the number of products on the Zero and One Sales prediction interval is also examined for OSPIs and TSPIs. Figure 15 (OSPIs) and Figure 16 (TSPIs) illustrate plots of the empirical Type I errors to show the effect of increasing product group sizes across four demand rate levels. The performance of the newly proposed prediction intervals are similar to the Zero Sales prediction intervals in the sense that the resulting empirical Type I errors are not near nominal values for groups of less than 200 products. In Figure 15, the empirical Type I error rates of the OSPIs are near nominal levels for groups involving more than 200 products and shorter MTBDs (MTBD = 100 and split group of MTBDs = 50/400 consisting mainly of products with MTBD = 50).

Figure 15. Empirical Type I error for zero and one sales OSPI across product group size.





Figure 16. Empirical Type I error for zero and one sales TSPI across product group size.



The new prediction interval generally did not produce acceptable empirical Type I error rates for MTBDs of 200 or 300 for the OSPIs and for an MTBD of 300 for the TSPI (except at the 99% confidence level). Figure 16 exhibits that the TSPIs for an MTBD of 300 typically has high empirical Type I errors, but they are mostly near the upper reliability boundary. Interestingly, for 200 products in Figure 16, the empirical Type I error rates for the Zero and One Sales TSPI cluster almost at the same point for each confidence level. This observation lends support to using 200 products as a reference point.

Figure 17 compares the performance of the Zero and One Sales OSPIs and TSPIs for a product group size of 200 across various MTBDs. At the 95% and 99% confidence levels, an MTBD of 10 yields a greatly inflated Type I error. For larger MTBDs, the empirical Type I error rates for the TSPI are lower than for the OSPI. The OSPIs are not reliable at the 99% confidence level except for at three MTBDs and even then the empirical Type I error rates are near the upper reliability boundary. To obtain a clearer comparison of Zero and One Sales OSPIs and TSPIs, the numeric values of the empirical Type I error rates are tabulated in Table 8. The bolded empirical Type I error rates indicate that the prediction interval is unreliable. For an MTBD of 250 or more, the OSPIs are not reliable. The empirical Type I error rate drops dramatically for the 95% and 99% confidence levels when the MTBD increases from 10 to 15. For MTBDs above 350, the 99% Zero and One Sales TSPI is not reliable. A separate analysis (not displayed) used 600 products, but failed to improve the reliability of the prediction intervals particularly for products with the same set of MTBDs above 200.

Figure 17. Empirical Type I error for zero and one sales OSPI and TSPI.




Table 8

Empirical Type I Error for Zero and One Sales OSPI and TSPI in <i>Figure 23</i>						
MTBD	One-Sided Zero and One Sales Prediction Intervals			Two-Sided Zero and One Sales Prediction Intervals		
	90%	95%	99%	90%	95%	99%
10	0.098	<b>0.090</b>	<b>0.075</b>	<b>0.109</b>	<b>0.098</b>	<b>0.051</b>
15	0.103	<b>0.063</b>	<b>0.019</b>	0.100	0.048	<b>0.016</b>
20	0.104	0.055	0.012	0.100	0.046	0.009
30	0.096	0.051	<b>0.014</b>	0.101	<b>0.058</b>	0.013
50	0.105	0.056	<b>0.014</b>	0.104	0.050	0.010
100	0.100	0.054	0.012	0.105	0.055	0.009
150	0.107	<b>0.058</b>	0.011	0.104	0.052	0.009
200	0.106	<b>0.059</b>	0.015	0.104	0.054	0.010
250	<b>0.110</b>	<b>0.062</b>	<b>0.015</b>	0.099	0.051	0.011
300	<b>0.114</b>	<b>0.064</b>	<b>0.018</b>	0.105	0.055	0.010
350	<b>0.113</b>	<b>0.062</b>	<b>0.018</b>	0.101	0.052	0.012
400	<b>0.124</b>	<b>0.066</b>	<b>0.020</b>	0.102	0.053	<b>0.016</b>
450	<b>0.124</b>	<b>0.072</b>	<b>0.023</b>	0.106	0.050	<b>0.014</b>
500	<b>0.128</b>	<b>0.072</b>	<b>0.022</b>	0.106	<b>0.059</b>	<b>0.014</b>
600	<b>0.122</b>	<b>0.069</b>	<b>0.026</b>	0.102	0.056	<b>0.019</b>
700	<b>0.129</b>	<b>0.073</b>	<b>0.027</b>	0.106	0.055	<b>0.018</b>
800	<b>0.139</b>	<b>0.080</b>	<b>0.026</b>	<b>0.109</b>	0.055	<b>0.022</b>
900	<b>0.114</b>	<b>0.075</b>	<b>0.028</b>	0.102	<b>0.057</b>	<b>0.016</b>
1,000	<b>0.137</b>	<b>0.071</b>	<b>0.023</b>	0.101	0.056	<b>0.024</b>

An important observation about Figures 13 through 17 is that the proposed Zero and One Sales prediction intervals perform more reliably with short MTBDs (higher demand rates). As mentioned previously, one possible explanation is that when the demand rate is higher, fewer products will have no sales, and more products will have periods with demand of one unit. Under these conditions, the Zero and One Sales prediction intervals are more reliable than the Zero Sales prediction intervals. If a Zero, One, and Two Sales prediction interval were developed as an extension for predicting the demand rate of products with no more than two sales, it seems reasonable to expect that this prediction interval would be reliable for products with a higher demand rate (shorter MTBD). This extension is not considered in this paper.

#### Research Question 4

How effective is the Bayesian approach to estimating optimal inventory levels for moderate-demand products as compared to one using a maximum likelihood estimator of the demand rate parameter of a Poisson process?

As described in Chapter 2, Popovic (1987) provides an estimation procedure for determining the optimal number of units to stock for each inventory product based on the costs of surplus and shortage. He uses a Bayesian approach in which a prior distribution must be established. This Bayes model will be compared with a traditional maximum likelihood approach, referred to as the Poisson model, in which the rate parameter of a Poisson distribution is estimated from the previous data. A Monte Carlo simulation of demand rates over ten time periods is performed to assess the effectiveness of each of these two methods.

The term, moderate-demand products, is used in this research question to mean products that are not intermittent. This research question is examined as a basis to examine the next two research questions.

A standard assumption often made about customer demand is that the sales of a product follow a Poisson distribution and each product's sale is independent of the sales of other products. Let the underlying unknown demand rate of a product be  $\lambda$ . The estimate of this rate could be zero for a product that has no demand or no recorded history of demand. Since the observed period of time may not be long enough to allow for an accurate prediction of future demand, a well-chosen prior distribution for the demand rate may provide a more reasonable forecast than a forecast of zero.

The Bayesian model may, for example, use a prior demand distribution for a product that is based on the distribution of demand for all products. In this approach, the parameters of the prior distribution are estimated using the entire pool of products. In this study, the gamma distribution is selected as a prior distribution for the demand rates of products since this distribution is mentioned as being an appropriate prior distribution for the  $\lambda$  parameter of the Poisson distribution (DuMouchel, 1999; Popovic, 1987; Hill, 1999). Some products may be selling fast and others may be selling slowly due to randomness in customer buying patterns. This approach is used to mitigate the effects of insufficient data in predicting future demand rates.

## Determining Optimal Inventory Levels Based on Costs and Distribution

Popovic (1987) shows that the optimal inventory level in the second unit time interval could be computed by considering the cost per time unit of a surplus  $C_1$  and a shortage  $C_2$  of an item as well as the demand distribution described by the a posteriori distribution. As additional demand information is accumulated over time, the a posteriori distribution of demand can be updated to improve the accuracy of the parameter estimates of the distribution.

Depending on the surplus cost and shortage cost of a single product, an inventory level,  $r_i^*$ , is determined at the end of the  $i^{\text{th}}$  time interval to minimize the cost. The optimal value  $r_i^*$  will satisfy one of the inequalities in (1) depending on whether the Bayes model or Poisson model is used. The present study examines both the performance of the Bayes model approach and the Poisson model approach in determining the optimal level of inventory using these inequalities.

Bayes Model Approach:

$$\sum_{k=0}^{r_i^*-1} \binom{\alpha+k-1}{k} \left(\frac{\beta}{\beta+1}\right)^\alpha \left(\frac{1}{\beta+1}\right)^k < \frac{C_2}{C_1+C_2} \leq \sum_{k=0}^{r_i^*} \binom{\alpha+k-1}{k} \left(\frac{\beta}{\beta+1}\right)^\alpha \left(\frac{1}{\beta+1}\right)^k \quad (1)$$

Poisson Model Approach:

$$\sum_{k=0}^{r_i^*-1} e^{-\lambda t} \frac{(\lambda t)^k}{k!} < \frac{C_2}{C_1+C_2} \leq \sum_{k=0}^{r_i^*} e^{-\lambda t} \frac{(\lambda t)^k}{k!}$$

Note that the form of the cost expression for inequality (1) requires that only the ratio of the surplus cost to shortage cost be known. That is, a shortage cost of \$5 and a surplus cost of \$1 will yield the same results as a shortage cost of \$50 and a surplus cost of \$10. A possible shortcoming of using these

approaches is that the costs for surplus and shortage may not be readily available for intermittent data. Parameter estimates should improve over time intervals as demand history accumulates.

A pool of 100 products is used in the simulation study to compare the costs based on the Bayes model and the Poisson model. Following Popovic (1987), demand rates for this pool of products are generated from a gamma distribution with a selected value of the mean equal to 3, which allows some products to be slow moving provided the standard deviation is not too small. Four standard deviations were selected: 5.5 (high), 1.7 (moderately high), 0.55 (low), and 0.17 (very low). Two parameters of the gamma distribution determine these values:  $\alpha$ , the shape parameter, and  $\beta$ , the scale parameter. The mean for the gamma distribution is equal to  $\frac{\alpha}{\beta}$  and the variance is equal to  $\frac{\alpha}{\beta^2}$ . The standard deviations resulted from the following four pairs of  $\alpha$  and  $\beta$  values:  $\alpha = 0.3$  and  $\beta = 0.1$ ,  $\alpha = 3$  and  $\beta = 1$ ,  $\alpha = 30$  and  $\beta = 10$ , and  $\alpha = 300$  and  $\beta = 100$ .

Once the demand rate  $\lambda$  is selected, actual product sales are generated using a Poisson process. The number of units sold for each product is simulated 10 times, representing sales over a period of 10 time units. This process is replicated 20 times and the average costs per time unit are recorded. The time units could be days, weeks, months or something else.

Three pairs of shortage costs and surplus costs selected for the simulation analysis are (0.1, 0.5), (0.5, 0.5), and (0.5, 0.1), so that the ratios are 1:5, 1:1, and 5:1. Therefore, four pairs of gamma parameters and three pairs of costs

parameters resulted in 12 simulation scenarios to evaluate the performance of the Bayes model and the Poisson model. The parameter values are varied in a series of experiments and three results are reported for each time period: theoretical minimum inventory costs, expected inventory costs for Bayes model, and expected inventory costs for Poisson model. The theoretical minimum inventory cost is computed by assuming that the demand rate and the demand distribution are precisely known.

As mentioned previously, this simulation study generates demand rates for 100 products using a gamma distribution, whose mean is set to the same value, but the standard deviation varies over four values. A generated demand rate is assigned to each product and product sales are simulated according to a Poisson process. Estimates of the demand rate are made using the Bayes and Poisson models after observing the generated sales. From these estimates, the amount of inventory to hold for the next time interval is computed based on the inequalities in (1).

Expected costs from holding this amount of inventory is then determined using the underlying distribution of demand for that product. The performances of the two models will be judged to determine which model has the lower expected cost. It follows that the expected costs for the Bayes model and for the Poisson model will always be greater than the computed theoretical minimum costs shown in each figure.

The results of the simulations addressing Research Question 4 are shown in Figure 18 through Figure 20. The three graphs have a mean of 3 and standard

deviation of 5.5. The surplus and shortage costs have been varied and are displayed above each graph. It is clear that varying the costs only changes the scale. The graphs also show that the models do not show a large difference in costs when the variance of demand rates of the products is larger. Initially, the Bayes model has lower costs but the difference is small. It can be observed that the costs drop substantially after the third time period in most cases for both models.

Figure 18. Cost per time period when demand mean is 3, standard deviation is 5.5 and surplus to shortage cost ratio is 1 to 5.




Figure 19. Cost per time period when demand mean is 3, standard deviation is 5.5 and surplus to shortage cost ratio is 1 to 1.




Figure 20. Cost per time period when demand mean is 3, standard deviation is 5.5 and surplus to shortage cost ratio is 5 to 1.




Figure 21 through Figure 23 reveal the results of the expected costs per time period when the standard deviation of the demand rates is moderately high. For these figures, the performances of the Bayes model and the Poisson model are fairly similar approximately after the fourth time period. In contrast to the previous three figures, there is a noticeable gap between the total costs for the Bayes model and the Poisson model for the first two time periods. As mentioned previously, the total cost per period using either model improves dramatically after only three periods. These three figures also demonstrate that there is an advantage to using the Bayes model particularly for the first couple of time periods when the variance of the product demand rates is moderately high. The ratio of the surplus cost to shortage cost affects the scale of the graphs, but the relative performance of the model does not change much.

Figure 21. Cost per time period when demand mean is 3, standard deviation is 1.7 and surplus to shortage cost ratio is 1 to 5.




Figure 22. Cost per time period when demand mean is 3, standard deviation is 1.7 and surplus to shortage cost ratio is 1 to 1.





Figure 23. Cost per time period when demand mean is 3, standard deviation is 1.7 and surplus to shortage cost ratio is 5 to 1.



Next, Figure 24 through Figure 26 show the results of the expected costs per time period when the standard deviation of the demand rates is low. For these figures, the performances of the Bayes model and the Poisson model are much different in the first few time periods. In contrast to the previous 6 figures, the total costs for the Bayes model are fairly close to the theoretically optimal total costs. Again, the Poisson model improves dramatically after several time periods, but takes longer for its total costs to be comparable to that of the Bayes model. These three figures demonstrate that there is an advantage to using the Bayes model particularly for the first five or six of time periods when the standard deviation of the product demand rates is low.

Figure 24. Cost per time period when demand mean is 3, standard deviation is 0.55 and surplus to shortage cost ratio is 1 to 5.




Figure 25. Cost per time period when demand mean is 3, standard deviation is 0.55 and surplus to shortage cost ratio is 1 to 1.





Figure 26. Cost per time period when demand mean is 3, standard deviation is 0.55 and surplus to shortage cost ratio is 5 to 1.



The next three figures reveal the results of the expected costs per time period when the standard deviation of the demand rates is very low. For these figures, the total costs of the Bayes model and the Poisson model are substantially different. Similar to the scenario when the standard deviation is low, the total costs for the Bayes model are very close to the theoretically optimal total costs and show little deviation. Although the Poisson model still improves dramatically after several time periods, the difference between the performances of the models is much greater. Figure 27 through Figure 29 demonstrate that there is a clear advantage to using the Bayes model when the standard deviation of the product demand rates is very low.

Figure 27. Cost per time period when demand mean is 3, standard deviation is 0.17 and surplus to shortage cost ratio is 1 to 5.




Figure 28. Cost per time period when demand mean is 3, standard deviation is 0.17 and surplus to shortage cost ratio is 1 to 1.





Figure 29. Cost per time period when demand mean is 3, standard deviation is 0.17 and surplus to shortage cost ratio is 5 to 1.



All 12 graphs addressing Research Question 4 suggest that as the number of time units increases, both the total costs based on the Poisson model and the Bayes model approach the theoretically optimal total costs after several time periods. This is understandable, since better estimates can be produced as more historical information is gathered. Furthermore, Figures 18 to 29 show that the ratio of surplus cost to shortage cost plays only a minor role in the relative performance of the two models. Clearly, the standard deviation of the demand rate of the products is the dominant factor in determining which model is better suited. When the standard deviation is large, little advantage is gained by using the Bayes model; however, as the products become more homogeneous, the Bayes model is superior to the Poisson model. Computationally, the Poisson model is easier to implement since no prior distribution needs to be selected. Only one mean demand rate was used in this simulation. Additional simulations are required to determine the effect of other mean demand rates on the performance of the models. In addition, other combinations of mean demand rate and standard deviation could alter the selection guidelines proposed for addressing Research Question 4.

#### Research Question 5

Research Question 5: How effective is the Bayesian approach to estimating optimal inventory levels for intermittent-demand products as compared to one using a maximum likelihood estimator of the demand rate parameter of a Poisson process?

In Research Question 4, the Bayesian approach to estimating optimal inventory levels relative to using the Poisson model was clearly affected by the standard deviation of the demand rate. The ratio of surplus cost to shortage cost had minimal affect on the relative performance of the two models. In this research question, the relationship between these models is more extensively analyzed for slow or intermittent demand.

The analysis for this question is similar to that presented for Research Question 4. For this study, intermittent demand will be regarded as having a mean demand rate of 0.1 or lower per time unit, which is the equivalent to an MTBD of 10 or larger. This rate was selected since in studying Research Question 1, an MTBD of less than 10 typically resulted in too few products having no sales in a unit of time for the proposed prediction intervals on the future demand rate to be reliable.

Figure 30 through Figure 32 show plots of the total cost per time period for each of the Bayes model and the Poisson model with the following demand rates in descending order for a pool of 100 products: 0.1, 0.05, 0.01, 0.005, 1/300, 0.002, and 0.001. For each of these demand rates, the  $\alpha$  and  $\beta$  parameters of the gamma distribution are selected so that for the standard deviation varies.

Table 3 summarizes the simulation results presented in the graphs with regard to which model – Bayes or Poisson – tends to have lower total costs over 10 time units. Because there are so few sales, the plots of the total costs are not always decreasing. For example, in most of the 19 figures both models show increasing trends over certain time periods.

Figure 30. Cost per time period when demand mean is 0.1, standard deviation is 0.0316 and surplus to shortage cost ratio is 1 to 5.




Figure 31. Cost per time period when demand mean is 0.1, standard deviation is 0.3162 and surplus to shortage cost ratio is 1 to 5.




Figure 32. Cost per time period when demand mean is 0.1, standard deviation is 1.0 and surplus to shortage cost ratio is 1 to 5.




Figure 33. Cost per time period when demand mean is 0.05, standard deviation is 0.0071 and surplus to shortage cost ratio is 1 to 5.




Figure 34. Cost per time period when demand mean is 0.05, standard deviation is 0.0224 and surplus to shortage cost ratio is 1 to 5.




Figure 35. Cost per time period when demand mean is 0.05, standard deviation is 0.0707 and surplus to shortage cost ratio is 1 to 5.




Figure 36. Cost per time period when demand mean is 0.05, standard deviation is 0.2236 and surplus to shortage cost ratio is 1 to 5.




Figure 37. Cost per time period when demand mean is 0.01, standard deviation is 0.0032 and surplus to shortage cost ratio is 1 to 5.




Figure 38. Cost per time period when demand mean is 0.01, standard deviation is 0.01 and surplus to shortage cost ratio is 1 to 5.




Figure 39. Cost per time period when demand mean is 0.01, standard deviation is 0.0316 and surplus to shortage cost ratio is 1 to 5.




Figure 40. Cost per time period when demand mean is 0.01, standard deviation is 0.1 and surplus to shortage cost ratio is 1 to 5.




Figure 41. Cost per time period when demand mean is 0.005, standard deviation is 0.0022 and surplus to shortage cost ratio is 1 to 5.




Figure 42. Cost per time period when demand mean is 0.005, standard deviation is 0.0071 and surplus to shortage cost ratio is 1 to 5.




Figure 43. Cost per time period when demand mean is 0.005, standard deviation is 0.0224 and surplus to shortage cost ratio is 1 to 5.




Figure 44. Cost per time period when demand mean is 1/300, standard deviation is 0.0005 and surplus to shortage cost ratio is 1 to 5.




Figure 45. Cost per time period when demand mean is 1/300, standard deviation is 0.0058 and surplus to shortage cost ratio is 1 to 5.




Figure 46. Cost per time period when demand mean is 0.002, standard deviation is 0.0014 and surplus to shortage cost ratio is 1 to 5.




Figure 47. Cost per time period when demand mean is 0.002, standard deviation is 0.0141 and surplus to shortage cost ratio is 1 to 5.




Figure 48. Cost per time period when demand mean is 0.001, standard deviation is 0.001 and surplus to shortage cost ratio is 1 to 5.




Figure 30 through Figure 48 exemplify the behavior of the proposed Bayes and Poisson models with various means and standard deviations for intermittent demand rates. Not surprisingly, the Bayes model tends to perform better in most cases, which is consistent with its performance in studying Research Question 4. However, the Poisson model performs well and produces lower costs in Figures 37, 38, 42, and 46. These observations indicate that the Poisson model tends to result in lower total inventory costs when the coefficient of variation is large. Simulations with larger coefficients of variations were conducted and indeed this general relationship appeared to hold.

In many of the 19 figures, there is a substantial increase in the cost for the Poisson model up to time period 5 and then a drop in time period 6. Even different random numbers generated by computer were used; however, this pattern still appeared to persist. Apparently, there are so few sales cumulatively

over the first five periods that the estimate of the demand rate is unstable. In many cases, the cumulative demand by time period 6 is enough to provide a stable estimate.

In general, Table 9 provides data that suggests that when the coefficient of variation is 200% or smaller, the Bayes model outperforms the Poisson model. Not all of the  $\alpha$  and  $\beta$  combinations listed in Table 9 are associated with displayed graphs. For the first seven combinations, graphs are not displayed to avoid an excessive number of graphs but their coefficients of variation are listed to illustrate the trend in the Poisson model doing well for larger coefficients of variation.

Table 9

Comparison of Bayes Model with Poisson Model Ranked by Coefficient of Variation (CV)

$\alpha$	$\beta$	$\mu$	$\sigma^2$	$\sigma$	Lower Cost Bayes or Poisson	CV
0.0001	0.01	0.010	1.000	1.00000	Equal - no inventory	10,000%
0.001	0.01	0.100	10.000	3.16228	Close	3,162%
0.001	1	0.001	0.001	0.03162	Equal - no inventory	3,162%
0.001	0.1	0.010	0.100	0.31623	Close	3,162%
0.002	1	0.002	0.002	0.04472	Bayes	2,236%
1/300	1	1/300	0.0033	0.05774	Poisson	1,741%
0.005	1	0.005	0.005	0.07071	Poisson	1,414%
0.01	1	0.010	0.010	0.10000	Poisson (Fig. 48)	1,000%

Table 10 Continued

$\alpha$	$\beta$	$\mu$	$\sigma^2$	$\sigma$	Lower Cost Bayes or Poisson	CV
0.01	0.1	0.100	1.000	1.00000	Poisson (Fig. 32)	1,000%
0.02	10	0.002	0.0002	0.01414	Bayes (Fig. 47)	707%
0.05	10	0.005	0.0005	0.02236	Bayes (Fig. 51)	447%
0.05	1	0.050	0.050	0.22361	Poisson (Fig. 36)	447%
0.1	10	0.010	0.001	0.03162	Bayes (Fig. 39)	316%
0.1	1	0.100	0.100	0.31623	Poisson (Fig. 31)	316%
1/3	100	1/300	0.00003	0.00580	Bayes (Fig. 45)	174%
0.5	100	0.005	0.00005	0.00707	Bayes (Fig. 42)	141%
0.5	10	0.050	0.005	0.07071	Bayes (Fig. 35)	141%
1	1,000	0.001	0.000001	0.00100	Bayes (Fig.48)	100%
1	100	0.010	0.0001	0.01000	Bayes (Fig.38)	100%
2	1,000	0.002	0.000002	0.00141	Bayes (Fig. 46)	71%
5	1,000	0.005	0.000005	0.00224	Bayes (Fig. 41)	45%
5	100	0.050	0.0005	0.02236	Bayes (Fig. 34)	45%
10	1,000	0.010	0.00001	0.00316	Bayes (Fig. 37)	32%
10	100	0.100	0.001	0.03162	Bayes (Fig. 30)	32%
1/300	10,000	0.00333	0.0000003	0.00058	Bayes (Fig. 44)	17%
50	1,000	0.050	0.00005	0.00707	Bayes (Fig. 33)	14%

For some very small means, i.e., Figures 44, 46 and 48, after a few time periods, the total inventory costs from the models tend to be close to the

theoretically optimal cost level. As seen in the study of Research Question 4, the total cost associated with the Bayes model is more cost efficient when the standard deviation of demand rates is smaller. It should be noted graphs were not included for all combinations of means and standard deviations investigated.

#### Research Question 6

How effective is the Bayesian approach using a mixture of prior distributions to estimate optimal inventory levels for intermittent-demand products as compared to one using a maximum likelihood estimator of the demand rate parameter of a Poisson process?

In Research Question 5, the Bayes model and Poisson model were evaluated over a range of intermittent demand rates. In this research question, a prior for the product demand rates after the first period is allowed to be a mixture of two priors. For the products that experienced two or more sales, the estimates for the demand rates for these products is the same as used in the Bayes model in the previous research question in which the gamma distribution was the prior distribution. Now, for the products experiencing no sale or one sale after the first period, the estimate for the demand rate is the prediction rate using the estimator  $\frac{M_1(t) + 2M_2(t)}{t}$  adapted from Ross (2002). This model using a mixture of the gamma distribution and the adapted Ross (2002) estimator and will be referred to as the Bayes/Ross model.

The Bayes/Ross model would be equivalent to the Bayes model used in studying Research Questions 4 and 5 if there were no products with fewer than two sales. Thus, the Bayes model and the Bayes/Ross model should be

comparable in performance except in those cases where there are many intermittent sales. To answer this research question, a simulation experiment similar to that conducted for addressing Research Questions 4 and 5 is carried out. One question of interest is: Will the Bayes/Ross model perform at least as well as the Poisson model under those simulation scenarios in which the latter was more cost-effective than the Bayes model?

A summary of the results based on the demand rates and standard deviations used in answering Research Question 5 is presented in Table 11. Figure 49 through Figure 55 are graphs displaying the performance of the Bayes/Ross model as compared to the Poisson model under several scenarios in which the coefficient of variation (CV) is large (mostly over 1000%) and the Poisson model was generally competitive. It is seen from Table 4, that the Poisson model's total cost was lower only in one case. It should be noted that even when the Bayes/Ross model is listed as having lower costs that does not mean it outperforms the Poisson model in every time period. For instance, in Figure 49, the Poisson model has lower total costs in time periods 3, 5, and 8. However, the initial trend favors the Bayes/Ross model and these are comparable in time periods 6, 7, and 9.

The results presented in Table 11 reveal that the Bayes/Ross model is a viable procedure for determining inventory levels for products with intermittent demand. Graphs for other scenarios are not included here since the pattern that the Bayes approach generally leads to lower costs has been established. Each alpha and beta combination considered in answering Research Question 5 and

presented in Table 9 was used in the simulations addressing Research Question 6 and are presented in Table 10. A comparison of Table 10 and 9 shows that the Poisson model outperforms the Bayes/Ross model less frequently than it does the Bayes model.

Table 10

Comparison of Bayes/Ross Model with Poisson Model Ranked by Coefficient of Variation (CV)

$\alpha$	$\beta$	$\mu$	$\sigma^2$	$\sigma$	Lower Cost Bayes/Ross or Poisson	CV
0.0001	0.01	0.010	1.000	1.00000	Same – no inventory	10,000%
0.001	0.01	0.100	10.000	3.16228	Bayes/Ross (Fig. 59)	3,162%
0.001	1	0.001	0.001	0.03162	Same – no inventory	3,162%
0.001	0.1	0.010	0.100	0.31623	Bayes/Ross (Fig. 60)	3,162%
0.002	1	0.002	0.002	0.04472	Bayes/Ross (Fig. 61)	2,236%
1/300	1	1/300	0.0033	0.05774	Bayes/Ross (Fig. 62)	1,740%
0.005	1	0.005	0.005	0.07071	Bayes/Ross (Fig. 63)	1,414%
0.01	1	0.010	0.010	0.10000	Close	1,000%
0.01	0.1	0.100	1.000	1.00000	Poisson	1,000%
0.02	10	0.002	0.0002	0.01414	Bayes/Ross	707%
0.05	10	0.005	0.0005	0.02236	Bayes/Ross	447%
0.05	1	0.050	0.050	0.22361	Bayes/Ross (Fig.64)	447%
0.1	10	0.010	0.001	0.03162	Bayes/Ross	316%
0.1	1	0.100	0.100	0.31623	Bayes/Ross (Fig. 65)	316%

1/3	100	1/300	0.00003	1.97699	Bayes/Ross	174%
0.5	100	0.005	0.00005	0.00707	Bayes/Ross	141%
0.5	10	0.050	0.005	0.07071	Bayes/Ross	141%
1	1,000	0.001	0.000001	0.00100	Close	100%
1	100	0.010	0.0001	0.01000	Bayes/Ross	100%
2	1,000	0.002	0.000002	0.00141	Bayes/Ross	71%
5	1,000	0.005	0.000005	0.00224	Bayes/Ross	45%
5	100	0.050	0.0005	0.02236	Bayes/Ross	45%
10	1,000	0.010	0.00001	0.00316	Bayes/Ross	32%
10	100	0.100	0.001	0.03162	Bayes/Ross	32 %
100/3	10,000	1/300	0.0000003	0.00058	Bayes/Ross	17%
50	1,000	0.050	0.00005	0.00707	Bayes/Ross	14%

---

To keep the number of displayed graphs from becoming excessive, seven plots are presented. The plots in Figures 49 to 55 show examples of alpha and beta combinations when the coefficient of variation was over 1,000% except for the two situations when the models produced no inventory. Additionally, two graphs when the Poisson model previously outperformed the Bayes model in Research Question 5 are shown. Most of the displayed plots are scenarios where the Bayes model did not clearly have consistent lower costs.

Figure 49. Bayes/Ross Model versus Poisson Model, demand mean is 0.1 and standard deviation is 3.162 with surplus to shortage cost ratio of 1 to 5.




Figure 50. Bayes/Ross Model versus Poisson Model, demand mean is 0.01 and standard deviation is 0.3162 with surplus to shortage cost ratio of 1 to 5.




Figure 51. Bayes/Ross Model versus Poisson Model, demand mean is 0.002 and standard deviation is 0.0447 with surplus to shortage cost ratio of 1 to 5.




Figure 52. Bayes/Ross Model versus Poisson Model, demand mean is 1/300 and standard deviation is 0.0577 with surplus to shortage cost ratio of 1 to 5.




Figure 53. Bayes/Ross Model versus Poisson Model, demand mean is 0.005 and standard deviation is 0.0707 with surplus to shortage cost ratio of 1 to 5.




Figure 54. Bayes/Ross Model versus Poisson Model, demand mean is 0.05 and standard deviation is 0.2236 with surplus to shortage cost ratio of 1 to 5.





Figure 55. Bayes/Ross Model versus Poisson Model, demand mean is 0.1 and standard deviation is 0.3162 with surplus to shortage cost ratio of 1 to 5.



### Research Question 7

Are the proposed prediction intervals investigated in Research Questions 1, 2, and 3 reliable for predicting future demand rates of slow-moving products using data from a major national retailer?

A database of sales from slow-moving products was obtained for 30 stores across all 50 states in the United States from a national retailer. Two separate analyses were performed. Each analysis used a different random set of sales data. The first analysis was intended to tabulate the number of reliable prediction intervals and display descriptive statistics. The second analysis was intended to gain insight into the effect that type of confidence interval, level of

confidence, and type of slow-moving product had on the reliability of the prediction interval.

In the first analysis, a random sample of 30 stores (10 small, 10 medium, and 10 large) was selected and the observed product sales were used to construct prediction intervals for the pool of products with zero sales and the pool of products with no more than one sale. The number of prediction intervals computed for this analysis was 360:30 stores times two types of product sales (products with zero sales and products with no more than one sale) times two types of prediction intervals (two-sided and one-sided) times three confidence levels (90%, 95%, and 99%). All 676 products common to the afore-mentioned 30 stores were used.

An observation period of 103 weeks (about 2 years of data) was split between a data set of observed product sales and a data set of future product sales. The prediction intervals for the future demand rate were estimated using the data set of observed product sales from either 12, 30, or 50 weeks. The reliability of the prediction intervals for future demand rate was assessed using the data set of future product sales from the remaining periods in the two-year data set. The lengths of the remaining number of weeks were 91, 73, or 53. From these weeks, the future demand rate of the pool of products having zero sales or no more than one sale was estimated.

Table 11 displays the number and percentage of stores with estimated future demand rate in the computed prediction intervals. The table is partitioned by type of prediction interval (one-sided or two-sided) and type of slow-moving

product (zero sales or no more than one sale denoted by Zero and One Sales). Table 12 shows the number and percentage of reliable prediction intervals across all the 360 possible scenarios.

As expected, proposed prediction intervals are not necessarily reliable for all situations when real data are used. The OSPIs contained estimated future demand rates in a high percentage of the cases with a pool of products exhibiting zero sales, but were reliable only about 60% of the time for a pool of products having zero or one sale. Note that for all 99% prediction intervals, the percentage of future demand rates inside the prediction intervals is larger than at the other two confidence levels. This occurs because the 99% prediction intervals are wider. The TSPIs for products with zero sales, performed poorly. Less than 10 percent of the stores had predicted demand rates inside of the prediction intervals. However, the Zero and One Sales TSPIs performed better with percentages of reliable intervals between 53% and 87%.

A question might arise as to how the TSPIs for products with zero sales can perform so poorly when the OSPIs for products with zero sales performs very well. The reason is because the future sales of several products during the remaining periods in the two year time frame were zero or close to zero. The OSPIs has a lower limit of zero, so these products do not affect its performance. Perhaps a longer time frame for the remaining periods would result in higher estimates for the demand rate of the products and improve the performance of the TSPIs.

Table 11

## Summary of Reliable Prediction Intervals across Retail Stores

One-Sided Prediction Interval for Products with Zero Sales for 30 Stores						
	Number and Percentage of Stores with Estimated Demand Rate Inside 90% OSPI		Number and Percentage of Stores with Estimated Demand Rate Inside 95% OSPI		Number and Percentage of Stores with Estimated Demand Rate Inside 99% OSPI	
12 Weeks of History	29	96.7%	29	96.7%	29	96.7%
30 Weeks of History	30	100.0%	30	100.0%	30	100.0%
50 Weeks of History	27	90.0%	28	93.3%	29	96.7%
One-Sided Prediction Interval for Products with Zero and One Sales for 30 Stores						
	Number and Percentage of Stores with Estimated Demand Rate Inside 90% OSPI		Number and Percentage of Stores with Estimated Demand Rate Inside 95% OSPI		Number and Percentage of Stores with Estimated Demand Rate Inside 99% OSPI	
12 Weeks of History	11	36.7%	16	53.3%	20	66.7%
30 Weeks of History	15	50.0%	20	66.7%	24	80.0%
50 Weeks of History	14	46.7%	16	53.3%	22	73.3%
Two-Sided Prediction Interval for Products with Zero Sales for 30 Stores						
	Number and Percentage of Stores with Estimated Demand Rate Inside 90% TSPI		Number and Percentage of Stores with Estimated Demand Rate Inside 95% TSPI		Number and Percentage of Stores with Estimated Demand Rate Inside 99% TSPI	
12 Weeks of History	2	6.7%	2	6.7%	3	10.0%
30 Weeks of History	0	0.0%	2	6.7%	3	10.0%
50 Weeks of History	5	16.7%	8	26.7%	14	46.7%
Two-Sided Prediction Interval for Products with Zero and One Sales for 30 Stores						
	Number and Percentage of Stores with Estimated Demand Rate Inside 90% TSPI		Number and Percentage of Stores with Estimated Demand Rate Inside 95% TSPI		Number and Percentage of Stores with Estimated Demand Rate Inside 99% TSPI	
12 Weeks of History	16	53.3%	18	60.0%	22	73.3%
30 Weeks of History	20	66.7%	22	73.3%	26	86.7%
50 Weeks of History	16	53.3%	18	60.0%	22	73.3%

Table 12

Numbers and Percentages of Reliable Prediction Intervals across 30 Stores

	Number of Reliable Prediction Intervals	Percentage of Reliable Prediction Intervals
12 Weeks of History	197	54.7%
30 Weeks of History	222	61.7%
50 Weeks of History	219	60.8%

The second analysis consisted of an experimental design to test the effect of type of confidence interval, level of confidence, and type of slow-moving product on the reliability of the prediction interval. Five stores were randomly selected to observe sales data over a period of 50 weeks. Two additional stores were also included with sales data observed over 12 weeks. These data were selected so as to allow adequate time to observe the sales of products, but with some sales data over a shorter period to evaluate the performance of the prediction intervals under conditions with few observed time periods.

For each of these seven samples, prediction intervals were computed for the slow-moving products across four families of products with identical part number prefixes and for the aggregate of the product families. This allowed for some homogeneity of sales rate among the products. The estimated future demand rate was compared to each proposed prediction interval. The number of future demand rates inside the prediction interval was recorded and is displayed in Table 13.

Table 13

Number of Future Demand Rates within each Interval across Product Families

1S_0or1_90	1S_0_90	1S_0or1_95	1S_0_95	1S_0or1_99	1S_0_99	2S_0or1_99	2S_0_99	2S_0or1_95	2S_0_95	2S_0or1_90	2S_0_90
2	3	2	3	3	3	3	3	2	2	2	2
3	5	4	5	4	5	4	3	3	1	3	1
1	2	1	3	2	3	2	3	2	2	1	2
1	4	1	4	4	5	4	3	3	1	1	1
2	4	2	4	4	4	4	1	4	0	2	0
2	4	3	4	3	4	4	2	2	2	2	2
0	5	5	5	5	5	5	1	5	1	5	1

Each computed prediction interval is determined by a combination of different levels of three factors: 90%, 95% or 99% confidence level, zero sales or no more than one sale, and one-sided or two-sided prediction interval. In Figure 68, the headings (as in 1S\_0or1\_90) use 0 or 0or1 to denote whether the prediction interval was computed for products with zero sales or no more than one sale. The 1S and 2S indicate one-sided and two-sided intervals, respectively, and the confidence level is indicated by the last two digits. To understand the effect of these three factors on the reliability of the prediction intervals, a three-way ANOVA was performed and the results are presented next in Table 14.

Table 14

## Three-Way ANOVA of the Factors Affecting Reliability of Prediction Intervals

Source	DF	Sum of Squares	Mean Square	F Value	Pr > F
Model	11	87.5595238	7.9599567	7.35	<.0001
Error	72	78.0000000	1.0833333		
Corrected Total	83	165.5595238			
Source	DF	Type III SS	Mean Square	F Value	Pr > F
One_or_Two_Sided	1	20.01190476	20.01190476	18.47	<.0001
CI_Level	2	19.73809524	9.86904762	9.11	0.0003
One_or_Two_*CI_Level	2	0.30952381	0.15476190	0.14	0.8671
SlowType	1	0.01190476	0.01190476	0.01	0.9168
One_or_Two_*SlowType	1	41.44047619	41.44047619	38.25	<.0001
CI_Level*SlowType	2	4.30952381	2.15476190	1.99	0.1443
One_or_Two_*CI_Level*SlowType	2	1.73809524	0.86904762	0.80	0.4523

The factor SlowType in Table 14 indicates whether the prediction of future demand rate was computed for products that had zero sales or no more than one sale. The factor CI\_Level took on one of the following three values: 90%, 95%, and 99%. The factor One\_or\_Two\_Sided had two levels: one-sided prediction interval or two-sided prediction interval.

The ANOVA illustrates that an interaction exists between the factor One or Two-Sided and the factor SlowType. The two main effects One-or-Two-Sided and CI\_Level are also significant at a reasonable significance level (i.e., 1%, 5%, or 10%). To further investigate differences across treatment combinations, two multiple comparison procedures – Tukey and Bonferroni - were selected because

they control the familywise error (FWE) rate (Kirk, 1995). The results are presented in Table 9. Note that factor level combinations with the same letter (either A or B) are not significant at the 5% significance level. Both procedures provided the same results.

Table 15  
Results from Tukey and Bonferroni Tests at 5% Significance Level


Grouping		Mean	N	Level
	A	4.1429	7	1S_99_0
	A	4.0000	7	1S_95_0
	A	3.8571	7	1S_90_0
	A	3.7143	7	2S_99_0or1
	A	3.5714	7	1S_99_0or1
B	A	3.0000	7	2S_95_0or1
B	A	2.5714	7	1S_95_0or1
B	A	2.2857	7	2S_99_0
B	A	2.2857	7	2S_90_0or1
B		1.5714	7	1S_90_0or1
B		1.2857	7	2S_95_0
B		1.2857	7	2S_90_0

There appear to be two clusters of factor levels that are significantly different: Cluster 1 with 1S\_99\_0, 1S\_95\_0, 1S\_90\_0, 2S\_99\_0or1, 1S\_99\_0or1 and Cluster 2 with 1S\_90\_0or1, 2S\_95\_0, and 2S\_90\_0. Cluster 1 includes all

the confidence levels for OSPIs predicting future demand for products with zero sales. This cluster also includes both the 99% OSPIs and the 99% TSPIs for products with Zero and One Sales. Cluster 2 includes two prediction intervals that are two-sided and used in predicting future demand rates for products with no sales. The other prediction interval in this cluster is one-sided but used for predicting products with no more than one sale at the 90% confidence level.

Another way to view the results is to plot the means of combinations of factor levels against the confidence level as illustrated in Figure 56. From this graph, one might conclude that the one-sided prediction interval for predicting future product demand with no sales is most reliable. In contrast, the two-sided prediction interval is least reliable.

Figure 56. Mean of factor levels versus confidence level.



## CHAPTER 5

### DISCUSSION AND IMPLICATIONS

This research addressed the problem of estimating future demand rates for products having a low demand rate. Three new methodologies were investigated in a simulation study across a variety of experimental conditions: product group size, mean time between demands (MTBD), Type I error levels, and group size combinations with MTBD. Each research question will be revisited with discussion and implications about the experimental conditions. This chapter was written so that a practitioner, unfamiliar with statistical methodology, may gain an understanding of the key contributions of the results. Guidance for the proper use of the proposed methods is provided. Previously used acronyms are spelled out.

#### Discussion of Research Question 1

Research Question 1: For intermittent-demand products, can reliable two-sided prediction intervals be derived using the  $\frac{M_1(t)}{t}$  estimator (Ross, 2002) to forecast the demand rate of products that have not sold over a specified time period?

This research question addressed the basic question of predicting the future demand rate of a group of aggregated products that have yet to experience a sale. In addressing this question, the phrase proposed prediction interval refers to two-sided prediction intervals (TSPIs) developed using the Ross

$\frac{M_1(t)}{t}$  estimator. An example of a TSPI for a demand rate could be an interval between 0.1 and 0.4 representing a range of values for future demand rates with some level of confidence such as a 95% confidence level. The lower and upper bounds of these proposed prediction intervals may be used by managers to make critical decisions about high cost slow-moving products. For example, a maximum and minimum cash flow can be more accurately estimated for high cost merchandise by using the lower and upper bounds on demand rate from the TSPIs and the price of the products. The cash flow information may provide retailers with essential sales information to make important decisions related to inventory management and operating costs.

For a particular retail store, a manager may only realize a small cost savings by knowing the bounds of a proposed prediction interval for estimating the demand rate for certain groups of slow-moving products. However, a company may have hundreds of stores selling these products and thus, efficiency on a grand scale may translate into a competitive advantage. A manager may decide to discontinue carrying a group of products once the demand rate falls below a certain economic threshold. In this case, a one-sided prediction interval may be more useful than a TSPI. This will be discussed later in Research Question 2.

The purpose of this simulation study was to determine conditions under which the proposed prediction intervals were reliable. What does reliable mean? When the nominally stated confidence level is close to the empirical confidence level, the prediction interval is considered reliable. A sales manager may interpret

reliability, in this instance, as the extent to which a prediction interval consistently includes the true demand rate over repeated sampling. If a 90% prediction interval is reliable, then over the long run 90% of the time the interval will contain the true demand rate.

Practitioners may be easily misled by the nominal confidence level of a proposed prediction interval while inappropriately applying the procedure, for example, to product group sizes that were too small. The proposed prediction intervals were referred to as being robust if they maintain their nominal confidence level under changing conditions. The proposed prediction intervals were robust for product group sizes of 200 or more. Furthermore, an implicit assumption was that product demand was observed for 100 time units. For this study, a time frame of 100 time units was selected as it provided sufficient time to observe demand using various underlying values for the demand rates in the simulation study. In addition, a time frame of 100 time units is a practical time period in which managers may need to assess seriously the viability of continuing to hold certain inventory. This time frame also approximated the total time units of data available from a Fortune 500 company.

Another parameter under which the proposed prediction intervals were examined for reliability was demand rate. The MTBD, which is the inverse of the demand rate, was varied to study the robustness of the proposed prediction intervals. If the MTBD was less than 20, that is, on average a sale occurs sooner than every 20 time periods, proposed prediction intervals were not reliable. In other words, if products were selling too fast, then the proposed prediction

intervals will not be meaningful. As shown in the simulation study, a proposed 90% prediction interval was reliable for MTBDs ranging between 30 and 800 time units. At the 95% and 99% confidence levels, the necessary MTBDs were between 20 and 800 time units and between 30 and 300 time units, respectively. Products with a very small demand, for example, an MTBD in excess of 800 time units, will be occurring too infrequently to provide enough information to construct proposed prediction intervals that are reliable.

A group of products may indeed have more than one distinct demand rate. For example, a group of footwear products might consist of various brands and types of styles. Shoes used by construction workers might have a distinctly different demand pattern than dress shoes. The proposed prediction intervals were examined under conditions of a mixture of two MTBDs in a product group. Their reliability is presented in Table 10. For brevity, only the 95% confidence level is illustrated. The letters R and U denote whether the proposed prediction interval is reliable (R) or unreliable (U), respectively. The columns of Table 17 represent the number of products in each subgroup that corresponds to the pairs of MTBDs. The two subgroups may consist of products that are dissimilar.

As Table 17 illustrates, very high or very low demand rates may affect the reliability. Unreliable combinations are bolded and reveal a pattern. Interestingly, a product subgroup with an MTBD of 10 time units (relatively fast moving) when paired with another product subgroup having an MTBD of 200, 400, or 1,000 time units was mostly unreliable, but when paired with a subgroup having an MTBD of 100 was reliable. If a high or low demand rate for one subgroup of products was

combined with another subgroup of products having demand rates that were favorable to the model's reliability, then acceptable results were obtained with respect to reliability.

Table 16  
Mixtures of MTBDs and Reliability for 95% Prediction Intervals

MTBD pair	Group size						
	25/175	50/150	75/125	100/100	125/75	150/50	175/25
100/10	R	R	R	R	R	R	R
200/10	U	U	U	R	U	R	R
400/10	U	U	U	R	U	U	R
1,000/10	U	U	U	U	U	U	U
10/50	R	R	R	R	R	R	R
100/50	R	R	R	R	R	R	R
200/50	R	R	R	R	R	R	R
400/50	R	R	R	R	R	R	R
1,000/50	R	R	R	R	R	R	R
200/100	R	R	R	R	R	R	U
400/100	R	R	R	R	R	R	R
1,000/100	R	R	R	R	R	R	R
200/400	U	R	R	R	R	R	R
1,000/400	R	U	U	U	U	U	R

This aspect of the study illustrated that combinations of demand rates may yield reliable prediction intervals. As a rule for practitioners, the demand rates of the products grouped together should not be a combination of a high demand rate and an extremely low demand rate, such as MTBDs of 1000/10 time units. Relative to the time that demand is observed, a subgroup of products with a similar MTBD provides sufficiently reliable TSPIs. That is, if demand is observed over 100 time units, then a subgroup of products with an MTBD of approximately

100 time units would provide reliable results. When possible, a group with a very small or large demand rate should be combined with a group with a more moderate demand rate to achieve reliable TSPIs.

#### Discussion of Research Question 2

Research Question 2: For intermittent-demand products, can reliable one-sided prediction intervals be derived using the  $\frac{M_1(t)}{t}$  estimator (Ross, 2002) to forecast the demand rate of products that have not sold over a specified time period to be incorporated in a stopping rule procedure?

Similar to Research Question 1, this question addressed the basic question of predicting the future demand rate of a group of aggregated products that have yet to experience a sale. However, in addressing this question, the phrase proposed prediction interval refers to one-sided prediction intervals (OSPIS) developed using the Ross  $\frac{M_1(t)}{t}$  estimator. As in Research Question 1, product groups do not need to be similar, but should have similar demand rates. Companies may have hundreds or even thousands of stock keeping units (SKUs) in their warehouse and must have a decision rule to reduce or discontinue slow-moving inventory. If the upper bound of a proposed prediction interval was below an economic threshold, then a manager will be compelled to discontinue a product line.

As illustrated in addressing Research Question 1, increasing product group sizes also improves the reliability of the proposed prediction intervals for Research Question 2. If the product group size was less than 200, the proposed

prediction intervals were not reliable with the exception of a product group having an MTBD of 100 time units and a confidence level of 90%. The simulation results revealed that the proposed prediction intervals for Research Question 1 were more reliable than those for Research Question 2. For example, for a product group size of 200, the proposed prediction intervals at the 90% level were reliable only between MTBDs of 30 and 200 time units, at the 95% level were reliable only between MTBDs of 50 and 100 time units, and at the 99% level were reliable only at an MTBD of 100 time units.

Why was it that the proposed prediction intervals close to an MTBD of 100 time units were reliable? One possible explanation was that product demand was observed for 100 time units and an MTBD greater than 100 time units does not result in enough demand in the given time frame to form a reliable prediction interval. For example, if products have an MTBD of 300 time units and the observed demand was over a time frame of 100 units, then many products will likely show zero or close to zero demand and the estimated demand rate will be too low.

If products have an MTBD of 20 time units under the same time frame, then by the end of 100 time units, few products will be left that have no observed demand. Unfortunately, managers often need to make decisions about slow-moving products over a time frame that is smaller than the MTBD of the products. Therefore, it was necessary to understand the range of MTBDs over which the proposed prediction intervals were reliable.

### Discussion of Research Question 3

Research Question 3: For intermittent-demand products, can reliable prediction intervals be developed using an extension of the estimators assessed in Research Questions 1 and 2 to forecast the demand rate of products that have sold no more than one unit over a specified time period?

This research question addressed the basic question of predicting the future demand rate of a group of aggregated products that have yet to experience a sale or that have experienced exactly one sale. Hence, the question was an extension of Research Questions 1 and 2. In addressing this question, the phrase proposed prediction interval refers to either the OSPIs or TSPIs for the demand rate of products that have no more than one sale over a specified time frame. Products with no more than one sale are a larger group than the products with no sales history and more likely to have a higher demand rate. If the proposed prediction intervals are determined to be robust over certain demand rates and product group sizes, then this may be extended to products with no more than two sales or with no more than three sales might be worth investigating. Thus, these results may be applied to additional slow-moving products that have sold no more than a few items.

To test this theory, a comparison of the reliability of the proposed prediction intervals (for products with no more than one sale) with the reliability of predictions intervals addressed in Research Questions 1 and 2 (for products experiencing zero) sales was conducted. The results revealed that the proposed prediction interval for this question tended to be more reliable for faster moving

products (smaller MTBD) and that the prediction intervals addressed in the previous two questions were more reliable for slower moving products. For example, for a product group of size 200, a 99% TSPI for products with no more than one sale was reliable between an MTBD of 20 time units and 300 time units whereas a 99% TSPI for products with no sales was reliable between 30 and 400 time units (although just barely at 400).

Table 17

Reliability of Prediction Intervals for Products of Sizes of 200 and 600

Product Group Size of 200 MTBD	One-Sided Prediction Intervals			Two-Sided Prediction Intervals		
	90%	95%	99%	90%	95%	99%
20	R	R	R	U	R	R
50	R	R	U	R	R	R
100	R	R	R	R	R	R
200	R	U	U	R	R	R
300	U	U	U	R	R	R
500	U	U	U	R	U	U
700	U	U	U	R	R	U
900	U	U	U	R	U	U

TABLE 17 Continued

Product Group Size of 600 MTBD	One-Sided Prediction Intervals			Two-Sided Prediction Intervals		
	90%	95%	99%	90%	95%	99%
20	U	U	U	R	U	U
50	R	R	R	R	R	R
100	R	R	R	R	R	R
200	U	U	U	R	R	U
300	U	U	U	U	R	R
500	U	U	U	R	R	R
700	U	U	U	U	U	U
900	U	U	U	U	R	U

As demonstrated in Research Questions 1 and 2, increasing the number of products resulted in a trend of improved reliabilities for the proposed prediction intervals. Since the product group size of 200 did not provide reliable prediction intervals for a number of MTBDs, especially large MTBDs, an additional study was conducted using a product group size of 600. Table 17 summarizes the results. An R and a U denote whether the proposed prediction interval was reliable (R) or unreliable (U), respectively. Interestingly, the larger product group size did not result in a substantial improvement in reliability. A possible explanation was that there was an MTBD level at which demand was too low to obtain sufficient data to provide reliable prediction intervals regardless of the size of the product group.

Similar to the results shown for Research Question 1, the proposed prediction intervals did not perform well for very high or very low demand rates. The 90% and 95% proposed TSPIs were found to be robust over a wide range of demand rates. Product group sizes much smaller than 200 are not displayed, but yielded prediction intervals that were more unreliable. The following guidelines are suggested for the proper use of prediction intervals.

1. Utilize product group sizes of at least 200 products.
2. Use prediction intervals for products having no more than one sale for relatively faster-moving products rather than the prediction intervals for products with no sales.
3. OSPIs are not reliable for MTBDs of 300 or more. TSPIs are more reliable for higher MTBDs but still are unreliable at various confidence levels for MTBDs of 300 or more. MTBDs less than 30 time units may also result in unreliable TSPIs or OSPIs.

#### Discussion of Research Questions 4 and 5

Research Question 4: How effective is the Bayesian approach to estimating optimal inventory levels for moderate-demand products as compared to one using a maximum likelihood estimator of the demand rate parameter of a Poisson process?

Research Question 5: How effective is the Bayesian approach to estimating optimal inventory levels for intermittent-demand products as compared to one using a maximum likelihood estimator of the demand rate parameter of a Poisson process?

Inventory managers periodically update their predictions of future demand rates for products. The Bayesian approach allowed for updates using historical data over specific time intervals. In this approach, the manager must assume a prior probability distribution for the demand rate of the products. This prior distribution may be based on a manager's experience with similar products. A variety of methodologies exists for helping the manager to determine a reasonable prior distribution. A survey of store managers may provide sufficient information about the likely demand of a product. A Delphi method (Dalkey, 1969) may be used to anonymously provide a panel consensus on the likelihood of sales. Thus, there may be some subjectivity in choosing a prior distribution.

In contrast to the Bayesian approach, a manager may assume that the demand for each product is Poisson distributed and estimate the demand rate parameter of this distribution. The manager may estimate the demand rate by dividing the number of historical sales by time. Using both of these approaches, this study examined the cost of ordering and holding inventory. In addition, since the true underlying demand rate was known in the simulation study, a theoretical optimal cost of inventory was computed and used as a benchmark.

The Bayesian approach in this study used a gamma distribution as a prior distribution of the demand rate. This approach is more involved than using a Poisson approach, but may be easily automated. In the simulation study, the expected cost of inventory for the Bayesian model and Poisson model were compared by varying the mean and standard deviation of the of the demand rates of 100 products, which may be functionally dissimilar. In addition, the cost

ratio of shortage and surplus inventory varied between 1:5, 1:1, and 5:1.

In answering Research Question 4, the mean demand was fixed at 3 and the standard deviations were varied from 0.17 to 5.5. The resulting total cost of the inventory using these approaches typically declined quickly over the first four or five periods and as time approached 10 periods, the last time period for which updates were computed, both approaches merged.

In answering Research Question 5, the mean demand varied over small values, namely from 0.001 to 0.1 and the standard deviations varied from 0.001 to 3.16. For these parameter values, little immediate improvement in inventory costs was noted using either approach. When the demand rate of the products varied greatly, that is, the standard deviation of the demand rates was large, the inventory cost using the Bayes model was not superior to that using the Poisson model. For products with homogeneous demand rates, the inventory cost using the Bayes model was lower than that of the Poisson model. Practitioners may still opt to use the Poisson model since a prior distribution for the demand does not need to be identified. According to simulation results, if the coefficient of variation of the demand rates of the products was 1,000% or more, then the Poisson approach will perform at least as well as the Bayesian approach.

#### Discussion of Research Question 6

Research Question 6: How effective is the Bayesian approach using a mixture of prior distributions to estimate optimal inventory levels for intermittent-demand products as compared to one using a maximum likelihood estimator of the demand rate parameter of a Poisson process?

This research question addressed a similar approach to that used in Research Question 5 except that the prior distribution was assigned differently for the Bayes model. For products experiencing two or more sales, the same prior distribution was used as in the previous question. For the remaining slow-moving products, the estimate for the demand rate was the Ross estimator used in the prediction intervals studied in Research Question 3. Hence, the prior distribution was considered a mixture of two prior distributions. Since the product variance substantially influenced the performance of the approaches used in answering the previous question, a Bayesian approach incorporating the Ross estimator was examined.

When all products have two or more sales in the first time period, this approach was identical to the Bayesian approach discussed in Research Questions 4 and 5, since the Ross estimator was not included in the computations. However, when a large percentage of the products, for example 30% or more, experience no demand, then this approach was appealing. The results for Research Question 5 revealed that the Poisson approach may be superior for small values of the coefficient of variations. For this question, the Bayes approach using the Ross estimator outperformed the Poisson model when the coefficient of variation was 7% or less.

Sales managers should consider the Bayes/Ross model as an alternative Bayesian approach when demand for products is intermittent but demand by other products are selling faster to minimize inventory costs. Estimates of the coefficient of variation may be used by managers in deciding which approach to

use. Simulation results suggested that the Bayes/Ross model may not lower inventory costs more than the Bayes model when the demand rates were homogenous across products.

#### Discussion of Research Question 7

Research Question 7: Are the proposed prediction intervals investigated in Research Questions 1, 2, and 3 reliable for predicting future demand rates of slow-moving products using data from a major national retailer?

Another way of stating this research question is: Does the Ross estimator work in a similar manner with actual data as it does with simulated data or does real data produce unexpected results? This research question examines the performance of the prediction intervals proposed in the first three research questions using a database of slow-moving products provided by a major national retailer. The products in this database were very slow moving, sometimes selling only one or two items in a year. With real data, there was no way of knowing the underlying distribution or the theoretical demand rate.

Despite limitations of knowledge about these rates, a random sample of 30 stores (10 small, 10 medium, and 10 large) was selected and observed product sales were used to construct prediction intervals. Sales data were available over approximately a 2-year period (about 100 weeks). Since a large number of observations were needed to estimate the future demand rate, the stores were observed over shorter time periods than the 100 time periods used in the simulation study. In practice, 100 time periods would have been easy to use for this company, but unfortunately, there would be insufficient future

observations to assess the prediction intervals' performance. Products were randomly selected so as not to select a group of products with a homogeneous demand rate.

Descriptive statistics on a sample of 30 stores revealed that the one-sided prediction intervals for products experiencing no sales were very reliable whereas the two-sided ones were not reliable. For product groups experiencing no more than one sale, the one-sided and two-sided prediction intervals' performances were similar with accuracy ranging from 36.7% to 86.7%.

Another analysis was conducted to test for the significance of any effect due to type of confidence interval (one-sided and two-sided), level of confidence (90%, 95%, and 99%), and type of slow-moving product (products either experiencing no sales or experiencing no more than one sale) on the reliability of the prediction intervals. Five stores were randomly selected to observe sales data over 50 weeks. ANOVA analysis identified an interaction effect to be significant between types of slow-moving product and whether the prediction interval was one or two sided as well as a significant effect due to confidence level. Further analysis of these effects resulted in the conclusion that for this data, the most reliable prediction intervals were one-sided prediction intervals for products with no sales and the least reliable were two-sided prediction intervals for product with no sales.

As a cautionary note, the results from analyzing these data were difficult to generalize since true future demand rates were not known as well as the degree of mixture of demand rates. However, the analysis provides insight on the

performance of the proposed prediction intervals using very slow-moving products. For practitioners using these proposed prediction intervals for slow-moving products, a longer period of time that approximates the MTBD of the prediction interval should prove more reliable.

### Managerial Implications

Inventory managers desire to stock products that customers want to purchase. This research addresses one aspect of this challenge: predicting the future demand rate of products that are slow moving, particularly products with no sale or with no more than one sale over a specified period of time. This section identifies guidelines for managers to use of the methods developed in this research.

To implement the proposed procedures in this research, two basic assumptions need to hold. Products are assumed to 1. be selling independently and 2. follow a Poisson process. What does this mean to the practitioner who wants to properly use the procedures? If products have a random demand pattern across time and there is no reason to believe sales over one interval of time will be different from another, then the Poisson process is a reasonable assumption. A statistical goodness-of-fit test should be performed to confirm that the number of SKUs sold is distributed as a Poisson distribution. The demand rate is assumed constant in a Poisson process. If products are selling because of cyclical promotions and advertising as well as seasonal effects, then the Poisson process will not be a good assumption and the procedures developed in this dissertation are not suitable. Independence of products is a basic assumption.

The sale of one product is assumed not to affect the sale of another product. Although these assumptions will never completely hold in reality, a manager may assess these assumptions by observing sales and determine if there are any outright violations.

Several questions were posed in this section that should assist inventory managers in implementing the proposed procedures. First, what are suitable values for  $n$  (number of products) and  $t$  (observed time) for constructing the proposed prediction intervals? A bound on the expected squared error for the prediction interval was shown to be  $\frac{1.07n}{t^2}$ . Thus, a bound on the margin of error for a prediction interval was  $Z_{\alpha/2} \sqrt{\frac{1.07n}{t^2}}$ . The manager may select the value of  $t$  as well as the value of the number of products,  $n$ , and the significance level  $\alpha$  that determines the  $100(1 - \alpha)\%$  confidence level to determine an acceptable margin of error for the future demand rate. A large enough  $t$  will make the prediction interval narrow.

Second, how should products be grouped when determining prediction intervals for products with no sales or no more than one sale? Management must decide which products are included in a group. The simulations revealed that as the product group size increases, the reliability of the proposed prediction intervals improved. This improvement was dramatic for product group sizes between 50 and 200 and more subtle for group sizes over 200.

Products may be grouped together by the similarity of their demand rate. These product groups may be products that are functionally dissimilar but that

are selling at similar rates. The simulation results demonstrate that a mixture of products at two distinct demand rates may result in reliable prediction intervals. For example, in electronics, the products comprising a group of resistors may vary by resistance levels from 1 ohm to thousands of ohms. Suppose that a mixture of two demand rates is present. For example, 1,000 ohms resistors experience moderate demand and the remaining resistors experience slow demand. Then the prediction interval for the demand rate of the mixture of the two types of resistors not selling would still be reliable under conditions studied in this simulation. Conditions that must be satisfied include a sufficient product group size, independence of purchases, and moderate demand rates.

Third, how slow do slow-moving products need to be? In Chapter 1, references were provided to research that investigated slow-moving products to illustrate demand rates that were considered slow moving. A definition of slow moving was presented in Chapter 1 as the average number of time periods between demands being 20% or more. This rate corresponds to an MTBD of 20 time units. This definition was motivated by the simulation study results since the prediction intervals were not reliable for smaller MTBDs (or equivalently faster demand rates). In the simulation study, slow was defined after examining the reliability of prediction intervals for MTBDs between 10 and 1,000 time units. Selected MTBDs ranged between 50 and 400 time units when investigating product group sizes in this study. A value of  $t$ , the time frame over which demand is observed, should be approximately near the MTBD for the product group. This is advised because the simulation results revealed that prediction intervals were

reliable across a wide variety of experimental conditions when an MTBD was specified as being 100 time units, the observed time frame selected for the study.

Third, when should one-sided and two-sided prediction intervals be used? One-sided prediction intervals are applicable to stopping rules to help determine when product demand rates are below threshold limits set by managers for carrying the merchandise. Knowledge of the upper endpoint of a one-sided prediction interval allows managers to compare this value to some threshold value for decision-making purposes. As long as estimated future demand rates are above an acceptable minimum determined by management, products will likely be kept in stock. Once the minimum demand rate (threshold value) is reached, products may be considered for liquidation. Two-sided prediction intervals are applicable to determining estimated demand rates for an entire family of products and may be used to determine the number of products to stock or to determine stock out rates. The OSPIs performed particularly well for the business sales data supplied by a national retailer. If estimated future demand is below an economically sound threshold value, inventory managers may conclude to discontinue the product.

Retailers that have a large inventory of products that have not sold may benefit from a rule to determine if the non-selling products should be replaced or liquidated at a reduced price. An example might be furniture sales. Suppose that 20 out of 100 sofa designs have not sold in Period 1. An assessment must be made as to the rate of sales for Period 2. If the estimated future demand rate for those sofa designs with no sales is below a certain economically feasible

threshold demand rate, then the 20 sofas designs should be discontinued.

Assume that the manager continued carrying all the sofa designs during Period 2. For example that at the end of Period 2, there are 15 sofa designs, possibly different from the previous 20 that did not sell, out of 100 total sofa designs that did not sell. A new estimate of the future demand rate for the subgroup of 15 designs is determined for Period 3. Again, a determination must be made as to whether to continue to hold these 15 designs.

Fourth, when should a Poisson model or a Bayes model be used? The results of this simulation demonstrate the viability of the Poisson model and the Bayes model for inventory decisions at the introductory point of a product's life cycle. Products such as the new 2007 Chrysler Aspen had no history of demand before September 2006. Although demand may be estimated from similar products, the actual demand remained to be determined. After the first period of sales (this may be a month), demand could be recorded separately for the various styles (or combination of options) of the Aspen. Either the Bayes model or the Poisson model may be used to estimate future inventory for this model at the end of each time period of sales. For styles that have similar demand rates, the Bayes model was demonstrated to clearly outperform the Poisson approach.

As an example of an application of the Bayesian and Poisson approaches, suppose that inventory decisions are important for minimizing costs in stocking the highly anticipated Sony Playstation 3 (PS3) in each of the various distribution channels or retail outlets. Depending on the expected demand variance of the products, the manager should select the appropriate model and compute

inventory levels at the end of certain time periods. After several periods, the models should show noticeable improvement in the total costs especially if the demand is not homogeneous across distribution channels.

Fifth, under what general conditions should the proposed prediction intervals and models be used? A manager is advised to use the proposed prediction intervals and models when actual circumstances mirror the conditions presented in the simulations. A general summary of the usefulness of the procedures and the conditions for using them are presented in Figures 63 and 64. An implicit assumption is independence of purchases. Oftentimes, customers purchase similar products and thus purchases are related. Thus, this assumption must be considered when using the proposed prediction intervals and Bayes model in applications to real-world data. The assumption of independence, while not holding precisely in reality, provides a simplifying assumption in assessing the performance of the proposed methodology.

*Figure 57. Usefulness of proposed prediction intervals and models.*

Two-sided prediction interval using Ross estimator	Used to predict future estimate of products with no sales or with only one sale. Need to record sales over a specified time period and then develop prediction intervals.
One-sided prediction interval using Ross estimator	Used when decisions must be made about continuing to hold inventory or discontinuing.
Bayes model	Use when demand for products is homogeneous or if the manager knows the historical distribution of the demand. This model may be used for fast- or moderate-demand inventory when demand forecasts are updated over specified time frames.
Poisson model	Used when there is high variance in the demand of products and demand forecasts are updated over specified time frames.
Bayes/Ross model	Used for very slow-moving products when demand forecasts are updated over specified time frames.

As noted in Figure 57, the number of products in a group and the demand rate of the group must be considered when using the proposed prediction intervals. When a substantially large group (at least 200 products in most cases) is available, reliable prediction intervals may be obtained for products with an actual MTBD of between 30 and 300 time units for either one-sided or two-sided prediction intervals. Relatively fast-moving products, that is, products with a short MTBD such as 10 or 20 time units and products that move slowly with longer MTBD such as 300 time units may produce misleading prediction intervals.

*Figure 58.* General conditions and assumptions for proposed prediction intervals and models.

TSPI using Ross estimator	A specified time frame must be selected to observe demand. A time frame of 100 time units is used for the results in this study. Appropriate MTBDs range from 30 to 400 time units. Advised number of products is at least 50, ideally 200+.
OSPI using Ross estimator	Assumptions are same as for TSPIs above. Appropriate MTBDs range from 30 to 200 time units. Appropriate number of products is at least 200. More products are needed than for TSPIs for same reliability.
Bayes model	Prior distribution is assumed to be gamma in this study. Alternative prior probabilities based on experience may be used. Appropriate when coefficient of variation is less than 1000%. Acceptable performance was observed for product group sizes of 100.
Poisson model	No prior distribution is needed. Appropriate when the coefficient of variation is greater than 1000%. Acceptable performance was observed for product group size of 100.
Bayes/Ross model	Prior distribution was mixture of gamma distribution and Ross estimator. Performance was acceptable over wide range for coefficient of variation. Acceptable performance was observed for product group sizes of 100.

The Bayes methodology is appropriate when a prior distribution of the demand rate is available and when surplus and shortage costs are known. Furthermore, the standard deviation of the products' demand rate will determine

if the Bayes or the Poisson model would be better suited. The Bayes model performed very well when groups of products were homogeneous with respect to their demand rates. When demand is known to be low and the coefficient of variation is high, as mentioned in the analysis using this approach, the maximum likelihood estimates for the Poisson process is advisable.

### Limitations

Simulations may be criticized for several reasons among which are an over simplification of real life conditions or unrealistic assumptions about relationships between data values. Despite running simulations over hundreds of conditions, only a limited number of specified parameter values were included in this research. Caution is advised when generalizing the results. Interpreting model performance outside of the parameters investigated may be misleading. Several limitations for the study are now listed.

- Simulations may not produce optimum solutions, but provide guidance on how a particular model will behave with a given set of inputs.
- Simulations are based on randomness due to their design. Despite many replications, interpreting borderline results may be difficult.
- Models used in simulations approximate reality and do not contain every possible relationship found in the real world.
- The demand for each product is considered independent of the demand of the other products. In reality, product demand is correlated with the demand for other products. The assumption of independence was a simplifying assumption.

Proposed methods under general correlation structures would make the problems addressed in this dissertation considerably more complex.

- Demand for products and sales of products were assumed to be equivalent. In practice, demand may occur, but may not result in a sale. This assumption was made to simplify the analyses.
- The Bayesian model and the Poisson model proposed in this research make several assumptions, which may not be applicable in practice. A Poisson process was assumed for the demand of goods. This was a general assumption often made in the literature. A gamma distribution was used for a prior distribution for the distribution of demand rates. There may be more appropriate priors for actual product sales. An assumption was made that Popovic's (1987) methodology was applicable to inventory with intermittent data. The values for the cost of surplus and shortages may be difficult to estimate.
- Parameter values selected in the Monte Carlo simulation study were limited to certain ranges. The interpretation of the results of the performance of the proposed models may be different over a more extensive set of parameters. For example, the value for the time frame over which sales are observed should be varied in relation to the MTBDs in the simulation study.
- Demand rates were assumed to be constant for products over time. External factors may change the demand rates. Pricing promotions or dramatic swings in the economy may have temporary effects on the demand of products.
- Due to computational constraints, 5,000 replications were used for each simulation. While this number was generally accepted as being large enough to

eliminate the effect of random extreme values, it was still possible that extreme values were randomly generated that were not representative.

- The national retail data used to confirm the analysis may not be comparable to that of other companies. Observed time periods and mixture of demand rates may differ.

#### Unique Contribution of Research and Future Research Ideas

Estimation of the future demand rate of a group of products without sales or with no more than one sale over a specified time period is difficult due to lack of data. There are limited demand rate estimation procedures for this type of slow-moving inventory. The proposed prediction intervals for the future demand rate of these slow-moving products are unique in that prediction intervals addressing this problem have not been presented. The Monte Carlo simulation results presented in this paper provide insight into conditions under which these prediction intervals are reliable.

Numerous research articles express the merits of a Bayesian approach to modeling inventory demand. However, the performance of these approaches using a Monte Carlo simulation study has not been performed over a variety of demand rates. Assessing the performance of a Bayesian approach to obtain optimal inventory levels for slow-moving inventory has not been previously presented and is a unique contribution to the literature on the merits of Bayesian approaches. The results in this research provide guidelines as to the effectiveness of a Bayesian approach with intermittent data.

Only a finite number of experimental conditions were investigated for the proposed methodology. Additional simulations should be completed to extend this research for values outside the ranges tested and even between the parameter values selected. For example, the proposed prediction intervals were studied over a specified range of product group sizes. General trends were identified, but running additional simulations with group sizes in between the points selected would support the general trend or possibly identify potential anomalies resulting from some particular group size.

Modified prediction intervals for future demand should be investigated to determine approaches to making them reliable to a wider range of demand rates and product group sizes. There may be a correction factor that may be developed to enhance the performance of the prediction intervals. In addition, prediction intervals for products having no more than two sales, or some given number of sales, should be developed and investigated.

The Bayesian approach to obtaining optimal inventory models should be investigated over a range of prior distributions and not just with a gamma distribution. Alternative multi-period inventory methodology for optimizing inventory levels should be explored. The Bayesian approach may be extended to more complex inventory management problems to account for inventory that has a limited life span or allows for replenishment in the middle of the single period.

Two limitations to the current study were the assumption of independence of the demand of products and the assumption that sales and demand were equivalent. Assuming a correlation structure for the demand of inventory

products would require newly proposed methodology that might be difficult to implement. Future research should address the issue in which independence of product demand does not hold and address the issue of estimating demand that may not result in a sale. These issues would require formulations that were more involved than those presented in this research.

Most importantly, future research should provide extensive guidelines to inventory managers to select reliable models to optimize inventory levels over a wide variety of product types. Furthermore, the performance of any proposed model must be interpreted so that inventory managers may use them appropriately. Future research should assess methodology that performs well under assumptions that mimic real world conditions.

## REFERENCES

- Abdel-ghaly, A. A., Chan, P. Y., & Littlewood, B. (1986). Evaluation of competing software reliability predictions. *IEEE Transactions on Software Engineering*, 12(9), 950-967.
- Archibald, B. D., & Silver, E. A. (1978). (s, S) policies under continuous review and discrete compound Poisson demands. *Management Science*, 24(9), 899-909.
- 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.
- Bagchi, U., Hayya, J. C., & Ord, J. K. (1983). The Hermite distribution as a model of demand during lead time for slow-moving items. *Decision Sciences*, 14(4), 447-466.
- Bartakke, M. N. (1981). A method of spare parts inventory planning. *Omega: The International Journal of Management Science*, 9(1), 51-58.
- Bartezzaghi, E., Verganti, R., & Zotteri, G. (1999a). A simulation framework for forecasting uncertain lumpy demand. *International Journal of Production Economics*, 59(1/3), 499-510.
- Bartezzaghi, E., Verganti, R., & Zotteri, G. (1999b). Measuring the impact of asymmetric demand distributions on inventories. *International Journal of Production Economics*, 60/61, 395-404.
- Box, G. E. P., & Jenkins, G. M. (1970). *Time series analysis: Forecasting and control*. San Francisco: Holden-Day.
- Boylan, J. E., & Syntetos, A. A. (2007). The accuracy of a modified Croston procedure. *International Journal of Production Economics*, 107(2), 511-517.
- Boylan, J. E., Syntetos, A. A., & Karakostas, G. C. (2007). Classification for forecasting and stock control: A case study. *Journal of the Operational Research Society*, doi: 10.1057/palgrave.jors.2602312.
- Brown, G. F. Jr., & Rogers, W. F. (1973). A Bayesian approach to demand estimation and inventory provisioning. *Naval Research Logistics Quarterly*, 20(4), 607-624.
- Burgin, T. A. (1972). Inventory control with normal demand and gamma lead times. *Operational Research Quarterly*, 23(1), 73-80.

- Burgin, T. A. (1975). The gamma distribution and inventory control. *Operational Research Quarterly*, 26(3), Part 1, 507-525.
- Burton, R. W., & Jaquette, S. C. (1973). The initial provisioning decision for insurance type items. *Naval Research Logistics Quarterly*, 20(1), 123-146.
- Celley, A. F., Clegg, W. H., Smith, A. W., & Vonderembse, M. A. (1986). Implementation of JIT in the United States. *Journal of Purchasing and Materials Management*, 22(4), 9-15.
- Chang, S.-L., Chung, K.-J., & Yang, W.-D. (2001). The (s, Q) inventory system for slow-moving items. *Journal of Information & Optimization Sciences*, 22(2), 383-395.
- Chopra, S., Reinhardt, G., & Dada, M. (2004). The effect of lead time uncertainty on safety stocks. *Decision Sciences*, 35(1), 1-24.
- Chung, K.-J., & Hou, K.-L. (2003). The determination of the optimal safety stock of the (s, Q) inventory system for slow moving items. *Journal of Statistics & Management Systems*, 6(2), 229-240.
- Chung, K.-J., & Huang, C.-K. (2003). A note on the (s, Q) inventory system for slow moving items. *Journal of Information & Optimization Sciences*, 24(3), 543-552.
- Cohen, M. A., Kleindorfer, P. R., Lee, H. L., & Pyke, D. F. (1992). Multi-item service constrained (s, S) policies for spare parts logistics systems. *Naval Research Logistics*, 39(4), 561-577.
- Cohen, M. A., & Lee, H. L. (1990). Out of touch with customer needs? Spare parts and after sales service. *Sloan Management Review*, 31(2), 55-66.
- Croston, J. D. (1972). Forecasting and stock control for intermittent demands. *Operational Research Quarterly*, 23(3), 289-303.
- Croston, J. D. (1974). Stock levels for slow-moving items. *Operational Research Quarterly*, 25(1), 123-130.
- Dalkey, N.C. (1969). *The Delphi method: An experimental study of group opinion*. Santa Monica, CA: RAND Corporation.
- Dekker, R., Kleijn, M. J., & Rooij, P. J. de (1998). A spare parts stocking policy based on equipment criticality. *International Journal of Production Economics*, 56/57, 69-77.

- De Wit, J. R. (1983). Inventory problems with slow moving items: A Bayesian approach. *The Statistician*, 32(1/2), 201-206.
- Dolgui, A., & Pashkevich, M. (2007). On the performance of binomial and beta-binomial models of demand forecasting for multiple slow-moving inventory items. *Computers and Operations Research*, in press.
- DuMouchel, W. (1999). Statistical practice: Bayesian data mining in large frequency tables, with an application to the FDA spontaneous reporting system. *The American Statistician*, 53(3), 177-202 (with discussions and replies).
- Dunsmuir, W. T. M., & Snyder, R. D. (1989). Control of inventories with intermittent demand. *European Journal of Operational Research*, 40(1), 16-21.
- Eaves, A. H. C., & Kingsman, B. G. (2004). Forecasting for the ordering and stock-holding of spare parts. *Journal of the Operational Research Society*, 55(4), 431-437.
- Ehrhardt, R. (1979). The power approximation for computing (s, S) inventory policies. *Management Science*, 25(8), 777-786.
- Feeney, G. J., & Sherbrooke, C. C. (1966). The (s - 1, s) inventory policy under compound Poisson demand. *Management Science*, 12(5), 391-411.
- Fisher, M. L., Hammond, J. H., Obermeyer, W. R., & Raman, A. (1994). Making supply meet demand in an uncertain world. *Harvard Business Review*, 72(3), 83-93.
- Forman, E. H., & Singpurwalla, N. D. (1977). An empirical stopping rule for debugging and testing computer software. *Journal of the American Statistical Association*, 72(360), 750-757.
- Fortuin, L., & Martin, H. (1999). Control of service parts. *International Journal of Operations and Production Management*, 19(9), 950-971.
- Freeland, J., & Porteus, E. (1980). Evaluating the effectiveness of a new method for computing approximately optimal (s, S) inventory policies. *Operations Research*, 28(2), 353-364.
- Gelders, L. F., & Groenweghe, P. F. (1985). Inventory models for slow moving items. *Belgian Journal of Operations Research, Statistics and Computer Science*, 25(1), 31-54.
- Gelders, L. F., & Van Looy, P. M. (1978). An inventory policy for slow and fast movers in a petrochemical plant: A case study. *Journal of the Operational Research Society*, 29(9), 867-874.

- Gentry, C. R. (2003). Riding fashion's roller coaster: Saks and New Look keep trend-setting items in stock. *Chain Store Age*, 79(6), 48, 50.
- George, J. A. (1982). Discounting a slow-moving stock item. *Interfaces*, 12(4), 53-56.
- Ghobbar, A. A. (2004). Forecasting intermittent demand for aircraft spare parts: A comparative evaluation of methods. *Journal of Aircraft*, 41(3), 665-673.
- Ghobbar, A. A., & Friend, C. H. (2002). Sources of intermittent demand for aircraft spare parts within airline operations. *Journal of Air Transport Management*, 8(4), 221-231.
- Ghobbar, A. A., & Friend, C. H. (2003). Evaluation of forecasting methods for intermittent parts demand in the field of aviation: A predictive model. *Computers and Operations Research*, 30(4), 2097-2114.
- Grahovac, J., & Chakravarty, A. (2001). Sharing and lateral transshipment of inventory in a supply chain with expensive low-demand items. *Management Science*, 47(4), 579-594.
- Haber, S. E., & Sitgreaves, R. (1970). A methodology for estimating expected usage of repair parts with application to parts with no usage history. *Naval Research Logistics Quarterly*, 17(4), 535-546.
- Haddock, J., Iyer, N. T., & Nagar, A. (1994). A heuristic for inventory management of slow-moving items. *Production Planning and Control*, 5(2), 165-174.
- Hadley, G., & Whitin, T. M. (1963). *Analysis of inventory systems*. Englewood Cliffs, NJ: Prentice-Hall.
- Harrington, L. H. (2003). Retail optimization: Time for change at Bulgari. *Transportation and Distribution*, 44(3), 54, 56-57.
- Harwell, M. R. (1991). Using randomization tests when errors are unequally correlated. *Computational Statistics and Data Analysis*, 11(1), 75-85.
- Heyvaert, A. C., & Hurt, A. (1956). Inventory management of slow-moving parts. *Operations Research*, 4(5), 572-580.
- Hill, R. M. (1999). Bayesian decision-making in inventory modeling. *IMA Journal of Mathematics Applied in Business & Industry*, 10(2), 147-163.

Hollier, R. H., Mak, K. L., & Lai, K. K. (2002). Computing optimal (s, S) policies for inventory systems with a cut-off transaction size and the option of joint replenishment. *International Journal of Production Research*, 40(14), 3375-3389.

Hollier, R. H., Mak, K. L., & Lam, C. L. (1995). Continuous review (s, S) policies for inventory systems incorporating a cut-off transaction size. *International Journal of Production Research*, 33(10), 2855-2865.

Horodowich, P. (1979). Evaluating the write-off of slow-moving finished goods inventory. *Management Accounting*, 60(9), 35-39.

Hua, Z. S., Zhang, B., Yang, J., & Tan, D. S. (2007). A new approach of forecasting intermittent demand for spare parts inventories in the process industries. *Journal of the Operational Research Society*, 58(1), 52-61.

Hummel, J. W., & Jesse, R. R. (1990). A spreadsheet heuristic approach for the stocking and retention of slow-moving, obsolescent items. *Computers and Industrial Engineering*, 18(2), 163-173.

Janssen, F. B. S. L. P., Heuts, R. M. J., & Kok, A. G. de (1998). On the (R, s, Q) inventory model when demand is modeled as a compound Bernoulli process. *European Journal of Operational Research*, 104(3), 423-436.

Johnson, E. (1968). On (s, S) policies. *Management Science*, 15(1), 80-101.

Johnston, F. R., & Boylan, J. E. (1996). Forecasting for items with intermittent demand. *Journal of the Operational Research Society*, 47(1), 113-121.

Johnston, F. R., Boylan, J. E., & Shale, E. A. (2003). An examination of the size of orders from customers, their characterisation, and the implications for inventory control of slow moving items. *Journal of the Operational Research Society*, 54(8), 833-837.

Kaufman, G. M. (1996). Successive sampling and software reliability. *Journal of Statistical Planning and Inference*, 49(3), 343-369.

Kennedy, W. J., Wayne Patterson, J., & Fredendall, L. D. (2002). An overview of recent literature on spare parts inventories. *International Journal of Production Economics*, 76(2), 201-215.

Kirk, R. E. (1995). *Experimental design: procedures for the behavioral sciences*. (3<sup>rd</sup> ed.). Pacific Grove, CA: Brooks/Cole.

Klein Haneveld, W. K., & Teunter, R. H. (1997). Optimal provisioning strategies for slow moving spare parts with small lead times. *Journal of the Operational Research Society*, 48(2), 184-194.

- Koulamas, C. (2006). A newsvendor problem with revenue sharing and channel coordination. *Decision Sciences*, 37(1), 91-100.
- Lee, H. L., & Nahmias, S. (1993). Single product, single location models. In S.C. Graves, A. H. G. Rinnooy Kan, & P. H. Zipkin (Eds.), *Handbooks of operations research and management science, Volume 4: Logistics of production and inventory*, (pp. 3-58). Amsterdam, the Netherlands: North-Holland.
- Levén, E., & Segerstedt, A. (2004). Inventory control with a modified Croston procedure and Erlang distribution. *International Journal of Production Economics*, 90(3), 361-367.
- Lyu, M. R. (Ed.). (1995). *Software reliability engineering*. Silver Spring, MD: IEEE Computer Society Press.
- Masters, J. M. (1993). Determination of near optimal stock levels for multi-echelon distribution inventories. *Journal of Business Logistics*, 14(2), 165-195.
- Matsuyama, K. (2006). The multi-period newsboy problem. *European Journal of Operational Research*, 171(1), 170-188.
- Miller, D. R. (1986). Exponential order statistic models of software reliability growth. *IEEE Transactions on Software Engineering*, 12(1), 12-24.
- Miragliotta, G., & Staudacher, A. P. (2004). Exploiting information sharing, stock management and capacity over sizing in the management of lumpy demand. *International Journal of Production Research*, 42(13), 2533-2554.
- Muckstadt, J. A., & Thomas, L. J. (1980). Are multi-echelon inventory methods worth implementing in systems with low-demand-rate items? *Management Science*, 26(5), 483-494.
- Omdahl, T. P. (1988). *Reliability, availability, and maintainability (RAM) dictionary*. Milwaukee: ASQC Quality Press.
- Popovic, J. B. (1987). Decision making on stock levels in cases of uncertain demand rate. *European Journal of Operational Research*, 32(2), 276-290.
- Price, B. A., & Haynsworth, H. C. (1986). How to prepare inventory forecasts for very low demand items. *The Journal of Business Forecasting*, 5(2), 21-22.
- Razi, M. A., & Tarn, J. M. (2003). An applied model for improving inventory management in ERP systems. *Logistics Information Management*, 16(2), 114-124.

- Regattieri, A., Gamberi, M., Gamberini, R., & Manzini, R. (2005). Managing lumpy demand for aircraft spare parts. *Journal of Air Transport Management*, 11(6), 426-431.
- Rosenfield, D. B. (1989). Disposal of excess inventory. *Operations Research*, 37(3), 404-409.
- Ross, S. M. (1985a). Statistical estimation of software reliability. *IEEE Transactions on Software Engineering*, SE11(5), 479-483.
- Ross, S. M. (1985b). Software reliability: The stopping rule problem. *IEEE Transactions on Software Engineering*, SE11(12), 1472-1476.
- Ross, S. M. (2002). *Introduction to probability models* (6<sup>th</sup> ed.). New York: Academic Press.
- Sani, B., & Kingsman, B. G. (1997). Selecting the best periodic inventory control and demand forecasting methods for low demand items. *Journal of the Operational Research Society*, 48(7), 700-713.
- Schneider, H. (1978). Methods for determining the reorder point of an (s, S) ordering policy when a service level is specified. *Journal of the Operational Research Society*, 29(12), 1181-1193.
- Schulmeyer, G., & McManus, J. (1999). *Handbook of software quality assurance* (3<sup>rd</sup> ed.). Upper Saddle River, NJ: Prentice-Hall.
- Schultz, C. R. (1987). Forecasting and inventory control for sporadic demand under periodic review. *Journal of the Operational Research Society*, 38(5), 453-458.
- Schultz, C. R. (1989). Replenishment delays for expensive slow-moving items. *Management Science*, 35(12), 1454-1462.
- Segerstedt, A. (1994). Inventory control with variation in lead times, especially when demand is intermittent. *International Journal of Production Economics*, 35(1/3), 365-372.
- Shale, E. A., Boylan, J. E., & Johnston, F. R. (2006). Forecasting for intermittent demand: The estimation of an unbiased average. *Journal of the Operational Research Society*, 57(5), 588-592.
- Shenstone, L., & Hyndman, R. J. (2005). Stochastic models underlying Croston's method for intermittent demand forecasting. *Journal of Forecasting*, 24(6), 389-402.

Silver, E. A. (1965). Bayesian determination of the reorder point of a slow moving item. *Operations Research*, 13(6), 989-997.

Silver, E. A. (1970). Some ideas related to the inventory control of items having erratic demand patterns. *Journal of the Canadian Operational Research Society*, 8(2), 87-100.

Silver, E. A. (1991). A graphical implementation aid for inventory management of slow-moving items. *Journal of the Operational Research Society*, 42(7), 605-608.

Silver, E. A., Pyke, D. F., & Peterson, R. (1998). *Inventory management and production planning and scheduling* (3<sup>rd</sup> ed.). New York: Wiley.

Smith, R., & Vemuganti, R. (1969). A learning model for inventory control of slow-moving items. *AIIE Transactions*, 1, 274-277.

Snyder, R. D. (1984). Inventory control with the gamma probability distribution. *European Journal of Operational Research*, 17(3), 373-381.

Snyder, R. D. (2002). Forecasting sales of slow and fast moving inventories. *European Journal of Operational Research*, 140(3), 684-699.

Stevenson, W. J. (2007). *Operations management* (9<sup>th</sup> ed.). St. Louis: McGraw-Hill/Irwin.

Strijbosch, L. W. G., Heuts, R. J. M., & Schoot, E. H. M. van der (2000). A combined forecast-inventory control procedure for spare parts. *Journal of the Operational Research Society*, 51(10), 1184-1192.

Syntetos, A. A., & Boylan, J. E. (2001). On the bias of intermittent demand estimates. *International Journal of Production Economics*, 71(1/3), 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. (2006). On the stock-control performance of intermittent demand estimators. *International Journal of Production Economics*, 103(1), 36-47.

Syntetos, A. A., Boylan, J. E., & Croston, J. D. (2005). On the categorization of demand patterns. *Journal of the Operational Research Society*, 56(5), 495-503.

Tekin, E., Gurler, U., & Berk, E. (2001). Age-based vs. stock level control policies for a perishable inventory system. *European Journal of Operational Research*, 134(2), 309-329.

- Teunter, T. H., & Vlachos, D. (2002). On the necessity of a disposal option for returned items that can be remanufactured. *International Journal of Production Economics*, 75(3), 257-266.
- Thompstone, R. M., & Silver, E. A. (1975). A coordinated inventory control system for compound Poisson demand and zero lead time. *International Journal of Production Research*, 13(6), 581-602.
- Tyworth, J. E., Guo, Y., & Ganeshan, R. (1996). Inventory control under gamma demand and random lead time. *Journal of Business Logistics*, 17(1), 291-304.
- Vereecke, A., & Verstraeten, P. (1994). An inventory management model for an inventory consisting of lumpy items, slow movers and fast movers. *International Journal of Production Economics*, 35(1/3), 379-389.
- Wagner, H. M., O'Hagan, M., & Lundh, B. (1965). An empirical study of exactly and approximately optimal inventory policies. *Management Science*, 11(7), 690-723.
- Ward, J. B. (1978). Determining reorder points when demand is lumpy. *Management Science*, 24(6), 623-632.
- Watson, R. B. (1987). The effects of demand-forecast fluctuations on customer service and inventory cost when demand is lumpy. *Journal of the Operational Research Society*, 38(1), 75-82.
- Weingart, Z. (1991). *Inventory models for slow-moving items for the Israeli Navy*. Unpublished masters thesis, Naval Postgraduate School, Monterey, CA, 98 pp.
- Whitin, T. M., & Youngs, J. W. T. (1955). A method for calculating optimal inventory levels and delivery time. *Naval Research Logistics Quarterly*, 2(3), 157-173.
- Wilcox, J. E. (1970). How to forecast 'lumpy' items. *Production and Inventory Management Journal*, 11(1), 51-54.
- 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.
- Williams, T. M. (1982). Reorder levels for lumpy demand. *Journal of the Operational Research Society*, 33(2), 185-189.

Williams, T. M. (1983). Tables of stock-outs with lumpy demand. *Journal of the Operational Research Society*, 34(5), 431-435.

Williams, T. M. (1984). Stock control with sporadic and slow-moving demand. *Journal of the Operational Research Society*, 35(10), 939-948.

Yeh, Q. J. (1997). A practical implementation of gamma distribution to the reordering decision of an inventory control problem. *Production and Inventory Management Journal*, 38(1), 51-57.

Zotteri, G. (2000). The impact of distributions of uncertain lumpy demand on inventories. *Production Planning & Control*, 11(1), 32-43.

Zwick, R. (1986). Rank and normal scores alternatives to Hotelling's  $T^2$ . *Multivariate Behavioral Research*, 21(2), 169-186.