


2012

Optimizing the Safety Stock Inventory Cost Under Target Service Level Constraints

Chetan T. Shivsharan
University of Massachusetts Amherst

Follow this and additional works at: <https://scholarworks.umass.edu/theses>

 Part of the [Industrial Engineering Commons](#), and the [Management Sciences and Quantitative Methods Commons](#)

Shivsharan, Chetan T., "Optimizing the Safety Stock Inventory Cost Under Target Service Level Constraints" (2012). *Masters Theses 1911 - February 2014*. 822.

Retrieved from <https://scholarworks.umass.edu/theses/822>

This thesis is brought to you for free and open access by ScholarWorks@UMass Amherst. It has been accepted for inclusion in Masters Theses 1911 - February 2014 by an authorized administrator of ScholarWorks@UMass Amherst. For more information, please contact scholarworks@library.umass.edu.

Optimizing the Safety Stock Inventory Cost Under Target Service Level Constraints

A Thesis Presented

by

Chetan Trimbak Shivsharan

Submitted to the Graduate School of the
University of Massachusetts Amherst in partial fulfillment
of the requirements for the degree of

MASTER OF SCIENCE IN INDUSTRIAL ENGINEERING AND
OPERATIONS RESEARCH

May 2012

Mechanical and Industrial Engineering

© Copyright by Chetan Trimbak Shivsharan 2012
All Rights Reserved

Optimizing the Safety Stock Cost Under Target Service Level Constraints

A Thesis Presented

by

Chetan T Shivsharan

Approved as to style and content by:

Ana Muriel, Chair

Hari Balasubramanian, Member

Ahmed Ghoniem, Member

Donald Fisher, Department Head
Department of Mechanical and Industrial
Engineering

ACKNOWLEDGMENTS

I thank my advisor Dr. Ana Muriel for her invaluable support, patience, guidance in the preparation of the thesis. Without her support and guidance this thesis would never have been possible. I want to thank Dr. Ahmed Ghoniem and Dr. Hari Balasubramanian for their helpful comments and suggestions during the completion of this thesis. I would also like to thank the Physics department who has supported me financially during my Master's program at University of Massachusetts Amherst. I also would like to specially appreciate my friends of the Supply Chain and Healthcare Management lab for all their help, support and advice.

Last but not the least I would like to thank my parents and my brother. Without their encouragement and support this work would not have been possible.

ABSTRACT

OPTIMIZING THE SAFETY STOCK UNDER SERVICE LEVEL

CONSTRAINTS

MAY 2012

CHETAN TRIMBAK SHIVSHARAN, B.E., UNIVERSTIY OF PUNE

M.S.I.E.O.R., UNIVERSITY OF MASSACHUSETTS AMHERST

Directed by: Professor Ana Muriel

The level of customer satisfaction largely depends on manufacturer's ability to respond to customer orders with promptness. The swiftness with which the manufacturers are able to meet customer demand is measured by the service level. There are two service level measures typically used. The first one is type 1 service level which denotes the probability of not stocking out over a planning period. The other is fill rate which denotes the proportion of demand satisfied with the existing inventory. We review the rich and diverse literature available on inventory cost optimization under these service level constraints. Subsequently two optimization models are developed for the two different types of service level measures. The goal is to determine the safety stock values for all products in a multi product inventory required to achieve aggregate type 1 and type 2 service levels at the minimum inventory cost. For both the models we also maintain a minimum threshold for individual type 1 and type 2 service level for every product. The models are solved using Lagrangian relaxation techniques.

The models are computationally solved in Microsoft Excel. We then carry out discrete event simulation to validate the results and to test the performance of the models. To provide the decision makers with an idea of variability in the service levels and the related risks associated with it on an immediate finite horizon planning scale we also carry out simulation for a time span of one, two and four years.

The results obtained show desired type 1 and type 2 service levels for products with under both infinite and finite planning horizons.

TABLE OF CONTENTS

	Page
ACKNOWLEDGMENTS.....	iv
ABSTRACT.....	v
LIST OF TABLES.....	x
LIST OF FIGURES.....	xii
CHAPTER	
1)INTRODUCTION.....	1
1.1 Background.....	1
1.2 Service Level Measurement.....	4
1.2.1 Type 1 service level	5
1.2.2 Type 2 Service Level / Fill rate.....	6
1.2.3 Other Service Level Measurements	7
2) LITERATURE REVIEW.....	15
3) PURPOSE OF WORK.....	34
4)ASSUMPTIONS AND NOTATIONS.....	35
4.1 Assumptions.....	35
4.2 Notations and Preliminary Calculations	36
5)MODELS.....	39
5.1 Model One for Type 1 Service Level Constraints	39

5.2 Model Two Type 2 Service Level Constraints or Fill Rate	40
6) SOLUTION TECHNIQUES.....	44
7) COMPUTATIONAL RESULTS.....	46
7.1 Data simulation	46
7.2 Model 1 Numerical Solution.....	47
7.3 Model 2 Numerical Solution.....	55
8) SIMULATION RESULTS.....	59
8.1 Simulation setup, assumptions and steps	59
8.2 Simulation results for 9 basic products.....	63
8.3 Simulation results for 27 classified products	72
8.4 Simulation results for normally distributed demand.....	85
9) RISK ANALYSIS.....	90
9.1 Simulation results for planning horizon of 6 years.....	91
9.2 Simulation results for planning horizon of 1, 2 and 4 years	97
10) LIMITATIONS.....	115
11) CONCLUSION.....	117

APPENDICES

A)LAGRANGIAN SOLUTIONS OF THE FIRST MODEL.....120

B)LAGRANGIAN SOLUTIONS OF THE SECOND MODEL.....123

C)ADDITIONAL GRAPHICAL RESULTS AND TABLES.....126

REFERENCES.....143

LIST OF TABLES

Table	Page
1: Sample data for service level calculations	5
2: Data simulation variables.....	46
3: Model 1 results summary.....	50
4: Model 2 results summary.....	56
5: Mean demand, lead times and cost of the 9 products for simulation.....	64
6: Example of Simulation carried out in Excel. Example is for Product 3 and for the first 10 weeks out of 10000 week planning horizon	65
7: Comparison of Type 1 service level and Fill rates obtained by simulation to expected solution.....	66
8: Standard deviations in Type 1 service level values and Fill rates for 20 replications of 9 products.....	71
9: Mean, Median, Skewness, Quartiles and Range of Type 1 service level values for 9 products across 20 replications.	72
10: Mean demand, lead times and cost for our experimental design.....	73
11: Variable information of the 27 classified products.....	74
12: Comparison of Type 1 service levels and Fill rates for the 27 classified products. The results obtained by simulation compared with the expected model solutions.	75
13: Discrete demand probabilities for slow moving demand.	88
14: Comparison of Type 1 service level and Fill rates obtained from simulations under poisson and normal distribution with the expected values from solution.....	88
15: Type 1 service levels over a planning horizon of 6 years. The average Type 1 service level and Fill rates for a 1000 trials with the minimum and maximum values.....	93
16: 95 percent confidence intervals, minimum and maximum Type 1 service level values for slow moving products for a planning horizon of 4 years.	102
17: 95 percent confidence intervals; minimum and maximum Fill rate values for slow moving products for a planning horizon of 4 years.	104

18: 95 percent confidence intervals; minimum and maximum Type 1 service level values for all products for a planning horizon of 2 years.....	106
19: 95 percent confidence intervals; minimum and maximum Type 1 service level values for all products for a planning horizon of 2 years.....	108
20: 95 percent confidence intervals; minimum and maximum Type 1 service level values for all products for a planning horizon of 1 years.....	111
21: 95 percent confidence intervals; minimum and maximum Fill rate values for all products for a planning horizon of 1 years.	113
22: Comparison of Type 1 service level and Fill rates obtained from simulations under poisson and normal distribution with the expected values from solution for products with medium to high demand.	134
23: 95 percent confidence intervals, minimum and maximum Type 1 service level values for medium and fast moving products for a planning horizon of 4 years.	141
24: 95 percent confidence intervals, minimum and maximum Fill rate values for medium and fast moving products for a planning horizon of 4 years.....	142

LIST OF FIGURES

Figure	Page
1: The number of products in different classes of Type 1 service level	51
2: Change in Type 1 service level values with change in mean demand of the products.....	52
3: Change in Type 1 service level values with change in cost of the products.....	52
4: The number of products in different classes of Fill rates	54
5: Service level comparison for nine basic products.	67
6: Fill rate comparisons for nine basic products.	67
7: Change in Type 1 service level and Fill rate with change in cost.	68
8: Change in Type 1 service level and Fill rate with change in lead times.....	69
9: Variability in results for the 20 simulation replications of Product 1 in Excel.	71
10: Type 1 service level comparison for 27 products.....	76
11: Fill rate comparisons for 27 products. Fill rates obtained by simulation compared with expected Fill rates solution.	77
12: Change in Service level, with change in mean demand and change in cost. (Low lead times).....	78
13: Change in Type 1 service level, with changes in demand for different lead times/cost product characteristics.	79
14: Change in Service level with change in lead time for different cost/demand product characteristics.....	80
15: Change in Service level with change in cost for different volume/lead time product characteristics.....	81
16: Change in Fill rate with change in the volume of demand and cost, for different lead times.	82
17: Changes in standard deviation of Type 1 service level values with changes in volume of demand and lead time respectively.....	83

18: Changes in standard deviation of Type 1 service level values with changes in cost	84
19: Changes in standard deviation of Fill rate with changes in cost.....	84
20: Comparison of Type 1 service level for slow moving products.....	89
21: Comparison of Fill rates for slow moving products.....	89
22: Variation in Type 1 service level for Product 3.....	94
23: Variation in Type 1 service level for product 6.....	95
24: Variation in Type 1 service level for Product 9.....	95
25: Variation in Fill rates for Product 3, Product 6 and Product 9.....	96
26: Type 1 service level variations for Products 1 and 3 for planning horizons of 1, 2 and 4 years.....	99
27: Type 1 service level variations for Products 6 and 9 for planning horizons of 1, 2 and 4 years.....	100
28: Change in Type 1 service level Quartiles for slow moving products for a planning horizon of 4 years.....	103
29: Change in Fill rate Quartiles for slow moving products for a planning horizon of 4 years.....	105
30: Change in Type 1 service level Quartiles for highlighted products in Table 18 for a planning horizon of 2 years.....	107
31: Change in Fill rate Quartiles for slow moving products (1-9) for a planning horizon of 2 years.....	109
32: Change in Fill rate Quartiles for highlighted products in Table 19 for a planning horizon of 2 years.....	110
33: Change in Type 1 service level Quartiles for slow moving products for a planning horizon of 1 year.....	112
34: Change in Type 1 service level Quartiles for highlighted products in Table 20 for a planning horizon of 1 year.....	112
35: Change in Fill rate Quartiles for slow moving products for a planning horizon of 1 year.....	114

36: Comparison of Type 1 service level and Fill rate for simulation of 27 products over infinite planning horizon.....	126
37: Change in Service level, with change in mean demand and change in cost. (Medium lead time).....	126
38: Change in Service level, with change in mean demand and change in cost. (High lead time)	127
39: Change in Service level, with change in mean demand and change in lead time. (Low cost)	127
40: Change in Service level, with change in mean demand and change in lead time. (Medium cost).....	128
41: Change in Service level, with change in mean demand and change in lead time. (High cost)	128
42: Change in Service level, with change in lead time and change in cost. (Low volume).....	129
43: Change in Service level, with change in lead time and change in cost. (Medium volume)	129
44 Change in Service level, with change in lead time and change in cost. (High volume).....	130
45 Change in Service level, with change in cost and change in volume. (Low Lead time)	130
46 Change in Service level, with change in cost and change in volume. (Medium Lead time)	131
47:Change in Service level, with change in cost and change in volume. (High Lead time).....	131
48: Changes in standard deviation of Fill rate values with changes in volume of demand.....	132
49: Changes in standard deviation of Fill rate values with changes in lead time.	132
50: Changes in standard deviation of Fill rate values with changes in cost.	133
51: Type 1 service level variations for Product 2 for planning horizons of 1, 2 and 4 years.	134

52: Type 1 service level variations for Product 4 for planning horizons of 1, 2 and 4 years.	134
53: Type 1 service level variations for Products 5 for planning horizons of 1, 2 and 4 years.	135
54: Type 1 service level variations for Product 7for planning horizons of 1, 2 and 4 years.	135
55: Type 1 service level variations for Product 8 for planning horizons of 1, 2 and 4 years.	136
56: Fill rate variations for Product 1 for planning horizons of 1, 2 and 4 years.	136
57: Fill rate variations for Product 2 for planning horizons of 1, 2 and 4 years.	137
58: Fill rate variations for Product 3 for planning horizons of 1, 2 and 4 years.	137
59: Fill rate variations for Product 4 for planning horizons of 1, 2 and 4 years.	138
60: Fill rate variations for Product 5 for planning horizons of 1, 2 and 4 years.	138
61: Fill rate variations for Product 6 for planning horizons of 1, 2 and 4 years.	139
62: Fill rate variations for Product 7 for planning horizons of 1, 2 and 4 years.	139
63: Fill rate variations for Product 8 for planning horizons of 1, 2 and 4 years.	140
64: Fill rate variations for Product 9 for planning horizons of 1, 2 and 4 years.	140

CHAPTER 1

INTRODUCTION

1.1 Background

The level of logistics service greatly influences customer satisfaction which in turn has a major impact on revenues (Ghiana, Laporte and Mussammo, 2004). In today's global market it is evident that companies offering superior customer service remain competitive and profitable (Larsen and Thornstenson, 2007). Competitive pressures in today's marketplace are forcing companies to offer quicker response to customer needs (Song, 2006). For nearly any business, an important aspect of customer service is the ability to responsively deliver products to customers (Ghiana, Laporte and Mussamo, 2004). As mentioned in the Kumar and Sharman's (1992) well-known essay "We love your product, but where is it?" on time delivery comes second after the product attributes, in deciding customer satisfaction. Meeting customer demand on time not only improves profit margins, but also develops a better public image. In an increasing competitive market hence it becomes important to meet customer demands on time.

Customer satisfaction or the ability to effectively respond to customer demand can be gauged by measuring service level (Nahmias 2007). Service level is defined in many ways; the simplest definition is the fraction of orders that are filled on or before their delivery due date (Nahmias 2007). It simply means having enough safety stock in the inventory so as to satisfy the customer demand. Holding an inventory can however be very expensive (Ghiana, Laporte and Mussamo, 2004). They state the following reasons,

- 1) A company that holds safety stock incurs an opportunity or capital cost represented by the return on investment the firm would have realized if the money had been better invested.
- 2) The warehousing costs must be incurred, which is made up of leasing cost and the operating and maintenance cost of the warehouse.
- 3) The company may also incur costs in form of insurance, shrinkage of the products or damage costs (Nahmias 2007).

Maintaining a low inventory cost is as important to a company as it is to achieve high service levels. The investment cost in the safety stock, along with a desire to maintain high level of service level, provides decision makers with a dilemma, which is difficult to deal with. In our current work we develop an inventory control system that achieves the desired aggregate service levels at a minimized cost. We develop two optimization models for the different types of service level measures. In both the inventory cost is minimized under each of the two service level measures. As there are two models, for the two most popular service level measures it is also necessary to compare the results obtained by them. We compare the models by looking at the difference in safety stocks, minimized cost and service levels obtained in each case.

After obtaining the solution to the models by Lagrangian relaxation we examine them further by carrying out discrete event simulation in Microsoft Excel. The simulation proves that the solution is driven by three variables, the volume of demand, lead time and the cost of each product. The solution is also indirectly

impacted by variability in demand. This impact is indirect we have assumed that demand follows Poisson distribution and hence is indirectly related to volume of demand. To gain more insight into the results we classify these three variables into low, medium and high and carry out further simulations. These simulations are carried in order to see how sensitive the service level measures are to change in the nature of the above variables.

There are many existing models and policies in the area of inventory cost optimization under service level constraints. These models though seem to work well only under specific assumptions and for a given set of conditions. One such supposition is assuming infinite planning horizon to achieve desired service level measures. The number of products in market that have continuous demand for a long amount of time is sparse. Managers and decision makers in industry are often looking at finite horizon planning and policies that have immediate consequences on their inventory level.

To ensure that the results obtained are implementable and robust on a finite horizon planning scale we finally carry out simulations for planning horizons ranging from one to six years.

The thesis is organized as follows. The thesis starts by investigating the different measures of service levels used in practice, clearly identifying the ones we will use in this research. This section is followed by a literature review of related previous work.

Chapter 3 describes the purpose and the specific aims of the project. Chapter 4 introduces the assumptions and mathematical notations used in developing the two models. The mathematical formulation of the two models is provided in Chapter 5. The next Chapter looks at Lagrangian solution techniques used to solve the models. We then look at results obtained for the two models and compare them in Chapter 6. Chapter 7 starts with describing the simulation model and subsequent computational results obtained by running simulations for different types of data sets. Chapter 8 provides results for the simulation of 27 classified products. The analysis and estimation of risk involved in implementing the model over a finite planning horizon is provided in Chapter 9. We discuss the limitations of our work in Chapter 10. The thesis concludes in Chapter 11 with a discussion of the results obtained, the conclusions derived and limitations of the work that lead to future research.

1.2 Service Level Measurement

In inventory theory the study of service level is as old as the theory of inventory itself (Silver 1970). Every text on operations or production management has sections that are devoted to address the problem of estimating service level. In the current section we take a look at a few of these important definitions. The most popular and frequently used definitions can be found in Nahmias (2007). There are two types of service levels defined as Type 1 and Type 2. The general idea of these definitions refers to the probability that demand or collection of demands is met from the inventory.

1.2.1 Type 1 service level

Type 1 service level is defined as the probability of not stocking out in the lead time (Wallace and Spearman, 2001). Denoted by α , the measurement of this type of service level is straightforward. α is defined as the proportion of cycles in which no stock out occurs (Nahmias 2007). Type 1 service level is used in cases where shortage occurrence has the same consequences independent of the time and the quantity involved. The following simple example illustrates the calculation of type 1 service level. We have different order cycles and the corresponding demand in each order cycle, with stock outs that have occurred in 2 of the order cycles.

Order Cycle	Demand in Cycle	Number of units that stocked out
1	180	0
2	240	0
3	200	0
4	350	125
5	500	0
6	250	50
7	155	0
8	375	0
9	275	0
10	190	0

Table 1 Sample data for service level calculations

In the table 1 there are ten cycles, out of which stock out happens in two cycles. The type 1 service level here is $\alpha = \frac{8}{10} = 0.8$ or 80 percent. In the current thesis there are two models the first of which is based on type 1 service level. Though type 1 service level maybe easy to compute, it is not very popular in industry. In type 1 service level the stock out for 1 unit is given the same weight as that given for stock out of a thousand units, thus making it a poor performance metric. In the first model, we will minimize the inventory subject to the type 1 service level.

1.2.2 Type 2 Service Level / Fill rate

Type 2 service level measures the proportion of demands that are met from the stock (Hopp and Spearman, 2001). This service level is denoted by β . β is the ratio of the total number of units that are satisfied with the existing inventory, to the total number of units that are ordered in a given time period (Hopp and Spearman, 2001). The type 2 service level is called as fill rate or line item fill rate. Type 2 service level will be hence referred to as fill rate, in the thesis. We can calculate the fill rate for the example in Table 1 above.

The total demand for the 10 order cycles is 2715, where there is a stock out of 175 units that happens in two order cycles. We are satisfying 2540 units of the 2715 units ordered, from the inventory. The fill rate in this case using the expression above is $\beta = \frac{2540}{2715} = 0.93$ or 93 percent whereas the type 1 service level was 80 percent. In the second model the inventory cost is minimized subject to a certain target fill rate.

1.2.3 Other Service Level Measurements

Among other important definitions for specifying service level is the work of Johnson et al (1995), in which they study the different existing definitions of fill rate. The various definitions of fill rate mentioned are for a periodic review (R, S) system. The inventory is reviewed every R units of time and is always brought up to a level S . They assume a replenishment lead time L which is constant and demands in different time periods are independent and identically distributed. The authors note that if for any order cycle i , there is stock out and if the replenishment order in period $i+1$ does not satisfy the previous stock out, then these unsatisfied stock outs are folded into the shortage period $i+1$ as well. This leads to the double counting of the units that are stocked out. As a result there is an overestimation of the actual number of units that are stocked out for a given cycle, which causes underestimation of fill rate.

The mathematical expression developed takes into account the number of units for a period's demand that are not filled. The authors measure $E(short)$, expected units short for a given period. While measuring $E(short)$ the expected number of stock outs, any backorders carried from the previous order cycle are subtracted, while any excess inventory left is added to the existing inventory. For any order cycle the number of units, that are stocked out in the corresponding order cycle are considered for calculating the fill rate. The expression developed works very well, for cases where a low fill rate is desirable. For cases of high fill rate the expression works as well as the existing

definitions of fill rate. In our current thesis we intend to achieve a high fill rate and hence will be using the definitions provided in the preceding section, as defined in the previous two sections.

Silver (1970) has also provided an alternate definition for type 2 service level or fill rate, to avoid double counting the units backordered. The author states that in case of small demand with large lead times there is double counting of units that are backordered. For example assume an inventory which is periodically reviewed every week. In week one the demand was 10 units out of which 8 of them were available in inventory. The fill rate is calculated using equation, $Fill\ rate = 1 - \frac{(Backorders)}{Demand}$ and for week one we obtain a value of 0.8. There is backorder of 2 units carried forward in week two's review. In week two the demand is 4 units and if only one is satisfied with the existing inventory there is a backorder of 3 units. The total number of unit's backordered in week two is 5. Using the same equation above now a negative fill rate is obtained. The expression of Silver (1970) is to account for this double counting of backorders. The authors first define Q as the order quantity received. The expression is based on the variable V . V is the fraction of the Q for a given period t and is the amount left after satisfying any backorders during the lead time. Depending on the reorder point OP and the demand during the lead time x the value of V will range from 0 to Q . If at time t the lead time demand x is less than the reorder point OP , then V takes the value Q . The entire new order received now is available to satisfy the demand at time period t . If lead time demand x is greater than reorder point OP but less than the sum of reorder point OP and order received Q , then V is equal to difference between the lead time demand and sum of reorder point and order

quantity. Finally if lead time demand x is greater than the sum of both the order quantity and the reorder point then the entire order is used to satisfy backorders with V being zero.

The fill rate is the ratio of partial expectation of the above three values of V to order Q .

$$\text{Fill rate} = \frac{\text{Expected value of } V}{Q}$$

where V is

$$V = \begin{cases} Q & \text{if } x \leq OP \\ OP + Q - x & \text{if } OP \leq x \leq OP + Q \\ 0 & \text{if } OP + Q \leq x \end{cases}$$

The above expression for Silver (1970) is for a periodic review inventory system. As we are using a continuous review inventory policy our first two models are based on the original type 1 and type 2 service level respectively. While carrying out the mathematical simulation we ensure that double counting of the backorders does not occur. This is done by first satisfying backorders for any given time period $t-1$ and then the leftover inventory is used to satisfied demand for time period t . The fill rate measurement is then done by taking the ratio of demand satisfied for period t , after accounting for backorders from the leftover inventory to the total demand realized for period t . Using the same example as above the fill rate after satisfying 2 backorders of the week one there is nothing left to satisfy the demand for week two. Hence the demand satisfied for week two is zero and the fill rate is zero.

In our work we use a $(S-1, S)$ continuous review policy. In such a policy when for every time period t a demand is fulfilled, the inventory level drops down below S and an order is placed to bring the inventory level back to S in the next $t+1$ period. Hence the order

that arrives in period $t+1$ can either be used to satisfy demand of period $t+1$ or the backorders of period t . This leads to a specific condition where for a continuous $(S-1, S)$ review policy the type 1 and type 2 service level measures are same. This result has been proved by Larsen et al (2007), and by Johnson et al (1995). These results by the mentioned authors imply that the optimal result for the two models should be similar or very close to each other.

Sobel (2004) in his work for achieving desired fill rates for a multistage supply chain has developed two fill rate approximations. The fill rate expression is for periodic review system that uses base stock policies. It is defined as β , which is the long run average fraction of the demand that can be satisfied immediately. The expression is based on convolutions for independent, identically distributed demand under normally distributed demand and demand with gamma distribution. There are two equations developed for β fill rate based on standardizing the demand and then double integrating this standardization. The approximations are computational simple and seem to account for double counting of backorders. In our current work however they do not seem to be accurate as there is a significant difference between the fill rates obtained by β equation and those obtained by simulation.

Zhang et al (2007) extend the above approximation to the specific case where inventory follows an periodic review (R,T) policy. This similar to (Q,r) policy where inventory is reviewed at periodic intervals of T and the inventory level is brought up to R .

In their paper on service level and spare parts, Cohen and Lee (1990), provide different definitions of fill rate that can be applied for measuring service level. The focus of their paper is customer manufacturer after product sales relationship. In particular they are dealing with spare parts which are required for the repairs of the products. The service level measurements described in the paper are;

- **Part Unit Fill Rate:** This measures the fraction of demand for a part shipped from on hand inventory in a given order cycle. The fill rate is specific for each part. The aggregate measure is the fraction of total demands for all parts shipped from on hand inventory.
- **Part Dollar Fill Rate:** This is a modification of the aggregate unit fill rate. Items are weighted by their value which is either their cost, the revenue they would generate or their contribution to the finished product.
- **Order Fill Rate:** Internal replenishment orders include requests for a wide range of items. These order fill rate measures the fraction of such orders that can be completely filled from on-hand inventory. This fill rate is relevant when the cost of shipping and receiving are high, or the fixed costs associated with the part unit Fill rate are high. In some cases, partial fill may not be allowed and an entire order is delayed because a small number of items are unavailable.
- **Repair Order Completion Rate:** This measure refers to the fraction of jobs not delayed by stock shortages. Here “job” refers to customer demand for service that leads to parts requirements. Since each job cannot be completed unless all the required parts are available, this measure is typically lower than the part fill rate

for an individual component. This rate captures service as experienced by the customer.

In literature a common method of measuring service level is by measuring backorders. In their famous METRIC (multi echelon technique for recoverable item control) paper Sheerbrooke (1968), measure service level by counting the number of unsatisfied orders. They aim to minimize these unsatisfied orders, defined as backorders. The work is done for the inventory of spare parts used in air plane repairs. The above work is extended by Brooks et al (1969) who uses four different ways to measure service level. Their work is also done for the Airforce base where air plane repairs are carried out. Among these four definitions two of them are backorders and type 2 service level or fill rate. They go on to define Operational Rate which means on any given day there are no backorders for air force repair depot. The authors have provided two methods for the measurement of Operational rate. The first is measuring the length of time in days where no backorders existed at the depot. The summation of these lengths is taken and then divided by 365 to provide operational rate. The second method is observing at the fixed time of a day whether backorders exist or not in the depot. The days with no backorders are summed and divided by 365 to give the operational rate. Their final measure of service level is NORS or non-operationally ready to supply aircrafts. This indicates the number of aircrafts which are still under repair present in the depot. Operational rate is calculated by adding the total number of days without backorders and then dividing the number by 365. NORS is calculated by adding the daily number of aircrafts under repair and then taking the average for a year.

Most of the work in service level measurements is for inventories of non perishable items. A few people have also tried to develop service level measures for perishable item inventories. Among the few people who have carried out research for perishable items the work of Donselaar et al (2010) is the one which provides alternate fill rate approximations. The fill rate equations are for (R,s,nQ) periodic review system. These approximations are based on calculating the “undershoot” defined as the difference between the inventory position and order up to level s , when a demand is generated.

Among other definitions for service level Inderfurth et al (1998) extend the definition of type 2 service level or fill rate to an entire planning period. The service level is defined as γ -type level. It is the ratio of expected backlog at the end of entire planning period to the expected demand for the entire period. The work of Inderfurth et al (1998) is done for a multi echelon system with the objective of minimizing the inventory cost subject to service level measures.

For multi echelon systems the authors Caggiano et al (2004) argue that the use of individual fill rates for measuring service level can be misleading. In their work on minimizing inventory cost across a supply chain, they measure service level using Channel fill rate. Channel fill rate is defined as the probability that incoming demand for a specific part at a specific location can be fulfilled by a specific period of time. Channel fill rate is required for cases when certain service level is to be achieved within a given time period.

Finally Caglar et al (2004) use response time as a measure of service level. Their work is carried out for inventory of items used in computer repairs. These repairs are carried out by a technician, who needs items to carry out the computer repairs. The response time is the time taken by the technicians to arrive at the customer site with a spare part to fix the machine after the customer reports a failure. In our present work we work with the first two definitions used for service level measurement and only for the case of the single echelon system.

CHAPTER 2

LITERATURE REVIEW

The literature on optimizing the inventory costs subject to service level constraints is extensive and can be divided into two parts. Most researchers have proposed cost minimization models while trying to achieve target service levels, whereas a few of them have looked to maximize the different service level measures, while dealing with a target inventory cost.

The earliest work dealing with measuring service level while addressing inventory control is the famous METRIC (multi echelon technique for recoverable item control) by Sherbrooke (1968). The work carried out at Rand Corporation is to improve the service level of spare parts that are used in the repair of aircrafts. The work is done for a two echelon supply chain network. The authors do not use type 1 service level or fill rate, but constrain the model developed by maintaining a specific limit on the number of backorders. The model minimizes the inventory cost subject to the constraint that the number of the backorders at a depot does not exceed a certain pre-defined value. As the inventory is made of spare parts that are used in the repair of the aircrafts, the demand of these parts is assumed to be sporadic. A continuous review ($S-I, S$) policy is followed, where the instant the demand falls below S , a part is ordered, to bring it back to S . Among the other important assumptions the demand is assumed to follow a Logarithmic Poisson distribution, a member of the Compound Poisson distribution family. The demand arrives in batches where batches are assumed to follow a Poisson distribution and then the

number of demands per batch has a logarithmic distribution. The cost minimization problem is solved with the help of Lagrangian multipliers and for the single item case. The results are then subsequently extended to the multiple item case. The model is implemented for the Hamilton Air force base and the George Air force Base.

Brooks et al (1969) have done related work to that of Sherbrooke, while trying to optimize the inventory of airport repair parts. The four different methods used by the authors to quantify service level are fill rate, backorders, operational rate and NORS (Not Operationally Ready Supply) aircraft for an airport base. Operational rate is defined as the probability that at any given date there are no back orders for the inventory base. It is measured by adding the number of days there are no backorders in a given year and then dividing this number by 365. The number of NORS aircrafts is calculated by observing the number of NORS aircrafts on a daily basis, and then taking the average for a year.

Same as the METRIC paper, a one to one ($S-I, S$) replenishment policy is observed with a demand that follows a compound Poisson process. The optimization model consists of minimizing the inventory cost, subject to each of the four different service level measures individually. The four different models are then tested, and the authors conclude that the model that aims to minimize the NORS aircrafts is more desirable as the results obtained are easier to implement operationally. The authors make use of the lagrangian multiplier to obtain solution to the model. The work of Aardal et al (1989) is similar to the work of Brooks et al (1969) and Sherbrooke (1968). The main purpose of the work is to understand the relation between service level constraints and the shortage costs. The

authors are trying to examine if the shadow price of a service level constraint can be interpreted as shortage cost. There are two cost minimization models for the two service level measures. Though the authors are able to relate service level constraints with shortage costs their work is done for a single item inventory following a periodic review (Q,r) policy.

The work of Graves et al (1982) is interesting and close to the idea of achieving a specified service level. The work is done for the individual inventory of a service representative, who carries this inventory to carry out repairs of large machines or equipment. If the inventory of the service representative lacks a particular spare item, which is needed for the large machines, then repair does not take place. It is also not practically feasible for the service representative to carry more than a specified number of spare parts. The model developed is an inventory cost minimization model, subject to a specific job completion rate α . α is type 1 service level and represents the combined probability for all spare parts being present. The decision variable x_i is a binary variable, which indicates whether the spare part is to be selected or not. The model has a knapsack like structure, and is solved with the help of a greedy algorithm where the spare parts are ranked in the order of preference till the constraint is satisfied.

Caglar et al (2004) also focus on the idea of achieving a specific service level. The work is for the items that are used in the repairs of computers. If the computer needs repairs it is because of some component which is malfunctioning. Either this spare item needs to be repaired or it needs to be replaced completely. There is a technician who carries out this

repair or replacement. The response time of the technician depends on whether the particular item is available in the inventory. The manufacturers try to have shortest possible response time, yet have a low inventory cost. The model developed by the authors has inventory cost minimized, subject to the constraint of having specific response times less than a specified threshold. The demand is assumed to follow a Poisson distribution and backorders are allowed. This work is done for a two echelon system made of a central warehouse and small depots where repairs are carried out. The inventory control policy is a base stock policy. In the constraints the response time is specified by measuring the waiting time W_j of the customer which is to be less than specific threshold. The waiting time is found out by using Little's law (Hopp and Spearman, 2001). The solution method is a heuristic which is based on Lagrangian decomposition of problems. The solution variable consists of finding S the base stock value at each warehouse and then deciding the same for all the depots. The numerical results show that the heuristic works very well for large problem instances. The recent work of Vliegen et al (2009) is very similar to the work of Caglar (2004). In this work inventory cost of tools used for repairs and servicing is optimized subject to fill rate constraints. An important issue the authors highlight is when service tools are demanded they are often required in combination with other tools. The demand of tools or parts is coupled. As the demand is coupled the fill rate of one tool impacts the fill rate of other products. Hence the authors feel it is important to achieve an order fill rate rather than an item fill rate. They develop three models based on Markov processes, to optimize the inventory cost subject to order fill rate constraints. In their results the performance of the three models is discussed.

The work of Schneider (1978) deals with finding the optimal reorder point s for an inventory model that follows an (S,s) ordering policy. The model is a single item model. The goal of the inventory control model is to achieve a desired service level. The two different service level measures used for computational analysis are type 1 service level and fill rate. There are two different solution methodologies used. The first one uses Brown's equation to find out backorders, while the other is based on asymptotic approximations. In both the solution methods the reorder point s is to be calculated. Though the author has presented mathematical formulation for both types of service levels the numerical results are computed only for type 2 service level or fill rate. The numerical results are carried out for both normal and gamma distributions of demand. The numerical results obtained by the asymptotic approximation are better than the ones obtained by Brown's equation. Schneider (1978) along with different researchers have extended the above work. The extended work focuses only on type 2 service level or fill rate. A different solution technique is also used to calculate s the reorder point. The work is to minimize the inventory cost subject to type 2 service level or fill rate measures. They define type 2 service level as, γ type service level. γ service level is a measure of expected backorders per order cycle.

The demand is assumed to be independent and identically distributed, and the inventory is monitored with a periodic review (S,s) policy. The lead time is assumed to be constant and the demand is assumed to be normally distributed. There are two solution methods, one based on Ehrhardt power approximation, while the other based on lagrangian

relaxation techniques. The power approximation is used to calculate variable D . D is the difference of S and s . The mathematical formulation is tested for a 288 item system. The total expected cost computed using the two different solution methodologies is compared. The results prove that Ehrhardt's power approximation provides results close to the optimal results.

Tijms et al (1984), have considered a general class of (S,s) models subject to type 2 service level constraints. Their work is an extension of work done by Schneider (1978) and others. The work is for inventory models where the periodic review may not be always possible. The authors feel that the work of Schneider (1978) is also difficult to implement in practice. The models developed are both for periodic and continuous review cases with the assumption of completely random lead times. For periodic review cases, the demand is assumed to follow a normal distribution and for the continuous review case it is assumed to follow a Poisson distribution. The lead time in both cases is assumed to be stochastic with a given mean and standard deviation. The number of cross orders is negligible. Among the other assumptions the most important assumption is the definition of Q , which is defined as the order quantity. Q is the difference between S and s , and is a predetermined value calculated using the standard EOQ expression (Economic Order Quantity equation). The solution methodologies include the use of two-moment's approximation for calculating s . The two-moment approximation provides 90 percent service level for periodic review models, and 95 percent service level for the continuous review models. Results are also obtained using the standard Lagrangian relaxation. The results validate the author's assumption of using the economic order quantity to find Q .

The results obtained by the two moment's approximation are close to optimal. In the case where the demand is very erratic, the reorder point s becomes very sensitive to the demand distribution making it difficult to achieve the high type 2 service level.

Cohen et al (1989) have considered spare parts inventory used for repairs. The stocking policy used by the warehouses is a periodic review base stock and order up to policy. The authors develop a generic inventory control model, which they go on to extend to different classes of customer. The other extension is for the case where commonality exists between different spare parts. The authors use fill rate or type 2 service level, as a measure of the service requirements. The work differs from others, as there are two constraints used to specify fill rate. The first constraint, defined as the chance constraint, is not a measure of fill rate for the spare parts, but for the fill rate of the entire product. The chance constraint specifies the lower limit that the fill rate should exceed. The other constraint is the part availability constraint, where the fill rate of each of the spare parts is now specified. The resulting problem is solved by taking the dual Lagrangian multiplier, and using a greedy heuristic. The model is extended to the case where customers are classified into low and high priority class and for use of common spare parts to carry out repairs. The computational results provide near optimal results for the two models and also for the case where commonality exists between the spare parts. Similar recent research is done by Samii et al (2011). where for different customers different fill rates are desired. A classification of customers like this leads to manufacturers holding certain amount of inventory for customers with high fill rate requirements. The authors define this process as nesting. This adds more complexity to the process of achieving desired fill

rate with a given inventory cost. The authors develop a fill rate equation which is applicable when inventory nesting is carried out. The solution to the cost minimization model provides near optimal results for the fill rate.

The work of Mitchell (1988) is for multi item inventory, with a desired Type 1 service level. The work done aims to achieve a specified type 1 service level for a single item, and at the same time trying to achieve a system wide service level for all the items. A periodic review system is used to control the inventory, with the (S, s) being the inventory policy. The lead time is assumed to be fixed, while demand is normally distributed. The non linear model is solved with the help of a generalized knapsack duality algorithm. The model is implemented for a 32 item inventory system. The data used to test the numerical results has a 20/80 structure, with 20 percent of the items accounting for 80 percent of the inventory cost. The generalized knapsack problem works well till a service level of 85 percent is achieved after which the solution procedure starts to degrade.

The work of Thonemann et al (2002) is for spare parts inventory and fill rate as the service level measure. The specific purpose of the authors is to show that an inventory system that tries to achieve an aggregate fill rate performs better than the one where the goal is to have specific individual item fill rate. The individual fill rate of the items has to be above a certain minimum fill rate value. At the same time the aggregate fill rate should be above a certain specific value. The model we have developed is trying to achieve similar results. The work differs as the focus is to calculate “skewness” which is the ratio

of fraction of different individual items demands to the total demand. The demand follows a Poisson distribution. The inventory follows a $(S-I, S)$ continuous review policy. The inventory cost is minimized subject to two constraints. For the first the individual item fill rate is above an agreed minimum value, while the other is a similar constraint for the entire order. There are two different models developed. In one case the expressions for cost and demand are exact and in the second case the same expressions are of a continuous nature. The exact model is solved with a simple heuristic, where the minimum base stock level for each part is chosen such that the fill rate is at least 50 percent. For each part the marginal increase in the demand weighted average fill rate is computed if the base stock level is increased by one. The part with the lowest marginal increase is chosen and the base stock level is increased by 1 by adding this part. The aggregate fill rate is calculated and the process is repeated till the desired aggregate system fill rate is achieved. The numerical results imply that in both the cases system based inventory policy has a significant lower cost even for very high fill rates.

The work of Song (1998) is an interesting contribution in the field of service level optimization. The work deals with the specific condition where there exists co-relation between demands of the different items. The service level measure though is type 2 service level or fill rate it is measured for a particular order and not for the entire demand in a given time period. The work is done for a continuous review base stock inventory which has to face Poisson demand. The solution method includes calculating order fill rate using one dimensional Poisson distributions. The mathematical analysis is done for two different cases of lead time. In one the author assumes equal lead times for all items,

while in other unequal lead times. The numerical results are for a two item inventory system with interdependent demand. The results conclude that the individual fill rate is not a precise indicator of the order fill rate. Also if the correlations of demand between the items change the proximity of the results with the optimal solution varies.

Bashyam et al (1998) have also carried out work which is a further extension of the work done by Schneider (1978) and that of Tjims et al (1984). In both the previous papers the difference between S and s is assumed to be very large. The approximations of Tjims et al (1984) perform poorly when the difference between S and s is small. The other important difference is the authors allow cross orders to exist. They assume that demands are not necessarily fulfilled in the same order in which they are received and yet a given level of service level has to be achieved. The service level measure is fill rate and follows a periodic review (S,s) policy. The model developed is a non linear model. The model is solved by developing local gradients or bounds on the cost function. The solution starts with finding Q using the economic order quantity. Subsequently the reorder point s is found out carrying a line search, after which a feasible directions search is employed as a check for feasibility. The work is one of the few that uses discrete event simulation to numerically test the model. There are different demand distributions observed to implement the model. The results indicate that as different steps of algorithm are carried out, the solutions get close to optimality. In most of the cases and demand distributions a fill rate of 90 percent is achieved. The authors mention that the computational effort is more for the latter stages of algorithm.

The work of Caggiano et al (2007) deals with achieving a desired service level within specific time intervals. The work is done for a multi echelon system. For different time intervals different fill rates are to be achieved. An example is fill rate of 90 percent is desired as soon as the order arrives, 8 hours after the order is received a fill rate of 95 percent is desired and finally after a few days from receiving the order the desired fill rate would be 98 percent. The authors intend to achieve not just a specific fill rate for the specified different time intervals, but also a desired Channel fill rate at the important nodes in different echelons. Channel Fill rate is the probability that an incoming demand for a specific part at a specific location can be fulfilled within a specific period of time. For the multi echelon network the authors feel that having a specific channel fill rate is more important than achieving a desired service level. The multi echelon network faces a Poisson demand and $(S-1, S)$ replenishment policy is used. The model assumes that backorders do exist. The objective function is made up of system inventory cost. The constraints are the different service levels to be achieved at different locations and for different time intervals. The aim is to find the base stock level S at each node. The three different solution methods used are Naïve Procedure, Fast Increment Procedure and a Primal Dual Lagrangian solution. The Primal Dual method carries out an iterative search, for the Lagrangian multiplier. In the Fast Increment Procedure and the Naïve Procedure the safety stock at higher echelons is first assigned, and then iteratively the safety stock at lower echelon points is updated until the desired fill rate is achieved. All the three solution procedures are tested for small, medium and large data sets. The naïve and fast increment procedures perform poorly for small problem sizes while the primal dual Lagrangian algorithm provides optimal results for all data sets, also while being fast.

The work of Sobel (2004) is aimed at achieving desired fill rate when using base stock policies. The author derives the equation for a single echelon system and then extends the work to multi echelon system. The fill rate equation is based on convolutions for demands that follow normal and gamma distribution. The results obtain provide base stock values required to achieve the desired fill rates for both single echelon and multi echelon supply chains. The fill rate equation is computationally simple though only applicable to base stock inventory control policies. Zhang et al (2007) and Kwon et al (2006) both have extended the work of Sobel (2004). Zhang et al (2007) extend it to a periodic review (R, T) policy, while Kwon et al (2006) propose a simulation based heuristic for serial inventory systems. In both cases the fill rate equation developed by Sobel (2004), is further approximated to meet the assumptions of specific cases.

Boyaci et al (2001) also focus on optimizing multi echelon inventory subject to service level measures. The work is for serial production or distribution systems. The service level measure used is fill rate. The work is for a continuous review system which uses a base stock policy, with the demand following a Poisson distribution. The first model is for a single echelon system, and then extended to a multi echelon case. The problem is solved by using two heuristics. In the first heuristic fill rate is calculated at each echelon. The echelons where a higher fill rate is needed are given priority and more safety stock is kept at these levels. These echelons are the end echelons that face the customer demand directly. In the second heuristic the system holds inventory only at specific levels. These echelons again are high fill rate echelons. The work is also extended to the case where fill

rate within a specific time window is to be achieved. The computational results provide near optimal results.

Among notable research done for multi echelon inventory subjected to different service levels, is the work of Ettl et al (2000). The service level measure is associated with calculating base stock levels at the different echelons. These base stock levels at different echelons are connected. The customers as well as the items are classified according to different service level requirements. The model is based on queuing theory equations which are used to calculate delays that happen on account of stockouts. The service level measure is fill rate and the demand follows a normal distribution. The model is developed for several purposes. Its first purpose is for performance analysis and the second for optimization. In case of performance analysis the inventory policies are first specified and then the corresponding fill rates obtained by using these inventory policies are calculated. The other purpose is optimizing the inventory cost subject to the service level requirements. The performance analysis is carried out modeling each site as an inventory queue.

The non linear model for optimization is solved with the conjugate gradient method. The model is implemented for 5 store network while carrying out performance analysis and optimization. For optimization the solution is a two step process where in the first stage 200 solutions are obtained and in then in the next step a conjugate search is again applied to the best of these 200 solutions. The numerical results are extended to the multi echelon case and in all cases near optimal results close to a fill rate of 95 percent are achieved.

Inderfurth et al (1998) also have carried out optimization of inventory cost for multi echelon system subject to service level measures. The focus of the work is on using safety stock buffers at different stages and also only for items that have a very high demand. Both type 1 and type 2 service level measures are used. The numerical results provide the different stocking points along with the base stock value S .

A few researchers have also investigated, optimizing service level subject to an inventory cost constraints. Schrady et al (1971) have developed a model where the time weighted shortages for inventory are minimized subject to the constraint that inventory cost does not exceed a certain value. The work assumes a (Q,r) inventory policy and demand is assumed to follow a normal distribution. The expression for time weighted shortage is the classical expression from Hadley and Whitin (1963) for a steady state probability distribution. There are two models where first one uses the classical expression for inventory cost while the second has simplified equations for inventory costs. The second model is solved with the standard Lagrangian relaxation. The first model is solved for an iterative search over the feasible region for variables (Q,r) . The model is numerically implemented for a 3 item inventory example and the authors feel it can be implemented to larger data sets.

Schroeder et al (1947) have developed multiple models for multiple measures of service levels. There are three different measures of service levels used. The first is expected number of actual backorders per year. The second is expected number of stock-out

occurrences per year. Finally the expected backorders at any point in time is the last measure of service level. In each of the three formulations the objective function seeks to minimize the above mentioned terms without exceeding the inventory cost beyond a certain limit. The inventory policy used is a (Q,r) policy. The demand and the lead time are normally distributed. All the three models are solved with the Lagrangian relaxation. The mathematical formulations are also extended to the multi item case. The mathematical solution proves that similar expressions are obtained for Q and r values for each of the three different models.

The paper of Schwarz et al (1985) deals specifically with fill rate optimization for one warehouse and an N identical retailer distribution system. The research is done for a single item case. The warehouse retailer supply network follows a (Q,r) policy. The lead time is assumed to be fixed and different nodes face Poisson demand. The objective function aims to maximize the fill rate which depends on the safety stock at the warehouse S_W and the safety stock at the retailer S_R . The only constraint is for the safety stock at the warehouse S_w and the total safety stock at all the retailers NS_R . The sum of these two safety stocks should not exceed a certain value S , which is the predetermined budget for the inventory cost. The solution involves calculating S_W and S_R the safety stock values at the retailer and the warehouse that will maximize the fill rate. The solution method consists of obtaining a policy and budget line on a graph of NS_R against S_w for different values of fill rate. The policy line or the fill rate line is the locus of all the points (NS_R, S_w) which are obtained as a solution to the optimization model for different values of budget S . The budget line is locus of points that satisfies the second constraint for

different budget values S . The intersection of these two lines provides the optimal safety stock values for retailers and the warehouse. The other solution methods include use of a Midband heuristic and the use of a Vertical heuristic. The computational results prove that both the heuristics provide results which are within one percent of optimality. The authors feel that the Vertical heuristic is easy to implement in practice.

Among other interesting research in this area is the work Chen et al (2001). The service level constraints used in most inventory models can be classified into two types. These two types are inventory models with minimal service level constraint and inventory models with mean service level constraint. The authors are trying to investigate the impact of difference in the service level constraints on the service levels achieved. They consider the two type 1 and type 2 service level measures. The work is done for a continuous (s,S) review policy. The work proves that inventory models with minimal service level models perform slightly better than those with mean service level constraints. Janssen et al (1999) consider the impact of data collection on type 2 service level or fill rate performance. The inventory follows a (R,s,Q) periodic review policy. The inventory is reviewed in intervals of R times, with s being the reorder point and Q the fixed order quantity. The demand process is described as compound renewal process a generalization of the compound Poisson process. Though the model is trying to optimize inventory cost subject to Fill rate constraints the focus of computational results is to find the optimal reorder point s .

Donselaar et al (2011) and Liu et al (1997) highlight the fact that not only products have infinite shelf life. Only a few products can be backordered while most of them may simply be perishable. The work of Donselaar et al (2011) develops a fill rate expression for a periodic review inventory system. The work of Liu et al (1997) classifies inventory into three parts: 1) products which can be completely backordered, 2) products which can be backordered if delivered within a certain time window and 3) products which can end up as lost sales. The model developed aims to optimize inventory cost of all the three types of products under fill rate constraints.

Type 1 and type 2 service level equations seem to be the most two widely accepted service level measures in industry. The work of Zheng et al (1999) evaluates the performance of two service level measures by comparing them in terms of cost, level of service and the inventory turnover ratio. The work though is done for a single stage continuous (s, Q) review policy. s is the reorder point, while Q is the fixed order quantity. After obtaining solution to the models we also carry out the comparison of our results in terms of safety stock obtained, the cost and the level of service level obtained.

Most of the research for inventory control under service level constraints is done for infinite planning horizon. As highlighted by Donselaar et al (2011), Samii et al (2011), not all products have an indefinite shelf life. Some of them perish quickly and may not be available for entire planning horizon. It is thus important to see how service level measures change when working under a finite horizon planning range. There are few people who have inspected this aspect of inventory control under service level constraint.

The most important is of Chen et al (2003) who prove that expected the finite horizon fill rate is higher than of infinite horizon fill rate. The above work is extended by Thomas (2005) who examines the above proposition under a base stock policy. Similar research is carried out by Ravichandra (2007) and Banerjee et al (2005) where both extend the above proposition to a multi echelon case under different types of inventory policies.

The intended work is very close to the work done by Hopp et al (1997), on simple inventory control policies for a manufacturer of mail processing equipment. The work is for a single echelon system, and the authors have proposed an optimization model, that minimizes the inventory cost, subject to satisfying a service level and with an additional limit on order frequency. The authors assume a (Q,r) policy and Poisson demand. The objective function is the average inventory holding cost; the system also incurs a backorder cost for items that are not met from the inventory. The model is solved with the help of Lagrangian multipliers. The authors have proposed three heuristics to solve the model. The first is based on the definition on type 1 service level and using the Lagrangian the optimal order quantity Q is decided. Subsequently the reorder point r is also decided. The second heuristic is based on the fill rate or type 2 service level. Again using the Lagrangian multipliers the quantities Q and r are calculated. The third heuristic developed is a hybrid heuristic where Q is calculated using the first heuristic, while r is calculated using the second heuristic.

To summarize the literature, though similar models have been developed by different people they seem to differ in terms of the solution approach. The service level measures

are case specific. As a result the focus of the work is only centered on a specific kind of service level measure, and for a particular kind of probability distribution for demand.

In our thesis we provide a general model for optimizing the inventory cost subject to two basic theoretical measures of service level. There are two models developed, one for the each of the two types of service level measures. The results obtained by both the models will then be compared. The models will be also tested for different types of demand distribution. The solution methodology is similar to the work carried out by Hopp et al (1997) however the model will be validated using a discrete event simulation approach. The underlying assumptions are also different as in the current work a $(S-1, S)$ policy is observed. The work of Hopp et al (1997) was for inventory based where demand was assumed to be Poisson distributed. In our work, we assume demand that follows normal distribution. We test the robustness of the models by assuming a demand that follows Poisson distribution in our results.

To ensure that our results are implementable we test the results obtained on a finite planning horizon along with infinite planning horizon.

CHAPTER 3

PURPOSE OF WORK

The purpose of our thesis is to develop an inventory control system for thousands of spare parts for a premier manufacturing system. The spare parts have a completely random and in some cases very sparse demand. As a result, achieving a target service level is difficult. To achieve the desired service level the number of spare parts in inventory would be high, resulting in a high inventory cost. The proposed work aims to strike a balance between inventory cost and the target service level. There are two service level measures defined in the preceding sections. The first model focuses on type 1 service level, that measures the frequency in which a stock out happens over a planning horizon, while the second model is for fill rate or type 2 service level, where we count the actual number of units that are met from stock. The two models optimize the safety stock costs. The difference in the resulting stock levels and costs associated with the two types of service levels will be compared. The solution of the models will provide the amount of safety stock to hold, to achieve the desired aggregate service levels under each of the constraints. The models will be solved with the help of Lagrangian relaxation, as was the case in most of the literature. We validate and test the robustness of the models carrying discrete event simulation in Microsoft Excel. The simulated inventory is made of products who have a huge variance in demand realized, lead times and cost. The simulation is carried out for both finite and infinite planning horizons.

CHAPTER 4

ASSUMPTIONS AND NOTATIONS

4.1 Assumptions

The objective function of both models is to minimize the inventory cost. The first model is constrained by type 1 service level, while the second is constrained by fill rate respectively. The aim of the first model is to achieve a specified aggregate type 1 service level for all the spare parts. For the second model we try to achieve aggregate type 2 service level or fill rate for all the parts. At the same time the two models are also trying to achieve a minimum type 1 service level and minimum fill rate for the individual parts.

The following are the key assumptions that are made;

- The models are multi product models, with thousands of parts involved in the process. The parts involved are spare parts used for repair purposes.
- The models are single echelon model with infinite planning horizon.
- The demand is assumed to be independent and identically distributed, with a mean of μ and a standard deviation of σ .
- The models are dealing with the spare items that are involved in the repair process and hence the items that are not available in the inventory can be backordered.
- The inventory review policy used is $(S-I, S)$ policy. The inventory level is kept constant at S units, if it falls below by one unit; an order is placed to bring the inventory level back to S .

- For the first model the type 1 service level is the aggregate service level. This is done to put more weight on parts that have higher demand while trying to achieve specified target type 1 service level. The aggregate service level is the ratio of the average demand to the total demand. A similar assumption is made for the second model where we are trying to achieve an aggregate target fill rate.
- The lead time is different for the spare parts and is assumed to be uncertain with a mean and standard deviation.

4.2 Notations and Preliminary Calculations

$Set I = \{1...N\}$; set of products from 1 to N.

c_i = inventory holding cost of product i .

μ_i = average demand of product i , for a single period

μ_{LTD}^i = average demand of product i , over lead time

σ_i = standard deviation of demand for product i

σ_{LTD}^i = standard deviation of demand for product i , over lead time

μ_{total} = average total demand of all the parts from 1 to N,

L_i = average lead time of product i

LS_i = standard deviation of lead time for product i

D_i = demand for product i , over lead time L_i , D_i is normally distributed with a mean of μ_{LTD}^i , and a standard deviation of σ_{LTD}^i , i.e. $D_i \sim N(\mu_{LTD}^i, \sigma_{LTD}^i)$.

z_i = standardized score for product i

$L(z_i)$ = standardized loss function associated with a particular z_i value.

MSL^i = Minimum type 1 service level required for individual spare part i . This is represented by z_i value in the standard normal distribution tables. The required minimum service level MSL is represented by $\Phi(z_i)$ which is cumulative distribution for standardized z values. This variable is applicable for the first model only.

MFR^i = Minimum type 2 service level or fill rate for the individual spare part i . This is represented or found out by looking up for z_i from the normal distribution table, for the desired fill rate that is to be achieved. The constraint associated with minimum fill rate involves the calculation of $L(z_i)$. This will be explained in the coming sections. This variable is applicable to the second model only.

d_i = average daily demand of product i , d_i is normally distributed with a mean of μ_i and a standard deviation of σ_i , i.e. $d_i \sim N(\mu_i, \sigma_i)$.

FR = desired Type 2 service level or fill rate for the second model.

SL_i = Type 1 service level for model for product i . This is not a minimum or maximum value rather a variable to represent actual type 1 service level associated with our decision variable z_i .

ϕ = probability distribution function of the standard normal distribution

Φ = cumulative distribution function of the standard normal distribution

S_i = order up to level for product i , for a time period.

fr_i = Type 2 service level or fill rate for any product i . This is not a minimum or maximum but a general variable to represent fill rate.

$fr_i(z_i)$ = Fill rate associated with order upto inventory policy. The order upto level is S_i for every product i , with a mean of μ_i and standard deviation of σ_i depending on the corresponding z_i value as explained in forthcoming sections. This also is not a minimum or maximum but a general variable to represent fill rate.

According to Nahmias (2005) text we have the following expression for average demand over lead time,

$$\mu_{LTD}^i = \mu_i * L_i \dots \dots \dots (1)$$

The lead time is different for different products. The safety stock calculations are to be done over a period of lead time. The safety stock equation for normally distributed demand with standardized score as given in Silver and Peterson (1979). As we are doing the calculations over a period of lead time, we take into account the standard deviation of demand over lead time,

$$Safety\ stock = z_i * \sigma_{LTD}^i \dots \dots \dots (2)$$

We assume that, lead time is uncertain and has a standard deviation. As lead time and demand can be considered as independent random variables using the expression from Silver and Peterson (1979) we calculate the standard deviation of demand over lead time as,

$$\sigma_{LTD}^i = \sqrt{(L_i * \sigma_i^2 + \mu_i^2 * LS_i^2)} \dots \dots \dots (3)$$

After calculating the standard deviation over lead time we then calculate the safety stock from equation 2 mentioned above.

CHAPTER 5

MODELS

The section is divided into two parts, each introducing the model associated with a particular service level.

5.1 Model One for Type 1 Service Level Constraints

$$\text{Minimize cost: } \sum_{i=1}^N c_i z_i \sigma_{LTD}^i$$

Subject to

$$\Phi(z_i) \geq MSL^i \text{ for } i = 1 \dots N$$

$$\sum_{i=1}^N \frac{\mu_i}{\mu_{total}} \Phi(z_i) \geq ASL$$

The objective function minimizes the total cost of the safety stock. The product of z_i and σ_{LTD}^i gives us safety stock value required to achieve the specified type 1 service level. The above model has two constraints. The first constraint specifies that type 1 service level for the individual spare parts should always be greater than a certain minimum value. The

decision variable in the model is z_i . The type 1 service level for each part i is represented by $\Phi(z_i)$ in the standardized normal distribution tables. It is required to be above a certain threshold MSL^i and hence constraints the minimum possible value for z_i . For example, assume we want to achieve minimum type 1 service level of 80 percent, we have $\Phi(z_i) = 0.7995$ in the standard normal distribution tables which means a corresponding z_i value of 0.84 is to be achieved.

The second and the more important of the constraints is to achieve a specified aggregate type 1 service level, ASL for all the products. The aggregate type 1 service level is the weighted average of the service level parts of each part. $SL_i = \Phi(z_i)$ where the weight is the ratio of the average demands μ_i of product i , to the total average demand of all products μ_{total} . The aggregate service level is a pre specified constant value.

5.2 Model Two Type 2 Service Level Constraints or Fill Rate

The above model can be extended to consider fill rate or type 2 service level measure. In this section we start off by deriving the exact mathematical expression for fill rate.

As defined in section 2, the fill rate definition is,

$$\text{Fill rate} = 1 - \frac{\text{expected number of backlogged units per cycle}}{\text{expected total number of units demanded per cycle}} \dots \dots \dots (4)$$

If fr_i is the fill rate in the above equation then for product i we have

$$fr_i = \frac{1 - E[\text{backorders}]}{\mu_i} \dots \dots \dots (5)$$

When the demand is normally distributed the standardized normal loss function is defined using Nahmias (2005) text.

$$L(z) = \int_z^\infty (t - z) \phi(t) dt = \int_z^\infty t \phi(t) dt - z(1 - \Phi(z)) = \phi(z) - z(1 - \Phi(z)). \quad (6)$$

The expected backorders for any time period for a given product i . can be calculated as follows

$$\frac{E[\text{Backorder}]}{\mu_i} = 1 - fr_i \dots \dots \dots (7)$$

$$\int_S^\infty (D_i - S_i) f(D_i) dD = (1 - fr_i) \mu_i \dots \dots (8)$$

Doing a change in variables, we have,

$$\frac{S_i - \mu_i}{\sigma_{LTD}^i} = z_i$$

$$S_i = \mu_i + z_i \sigma_{LTD}^i \dots \dots \dots (9)$$

Substituting and carrying out similar substitutions given in Silver, Peterson (1979) text;

$$\phi(z_i) - z_i(1 - \Phi(z_i)) = \frac{\mu_i}{\sigma_{LTD}^i} (1 - fr_i)$$

Substitute as $L(z_i)$ from equation 6 we get,

$$L(z_i) = (1 - fr_i) \frac{\mu_i}{\sigma_{LTD}^i} \dots \dots \dots (10)$$

As fill rate fr_i is now associated with an corresponding order upto level S_i with a mean of μ_i and a standard deviation of σ_{LTD}^i for every product i , we replace it with $fr_i(z_i)$.

Rewriting above equation for fill rate,

$$fr_i(z_i) = \left(1 - \frac{L(z_i) \sigma_{LTD}^i}{\mu_i} \right) \dots \dots \dots (11)$$

The fill rate for all the individual items is to be always kept above a minimum value MFR^i . MFR^i is a given constant for all the products. Hence equation 11 can be rewritten as;

$$fr_i(z_i) = \left(1 - \frac{L(z_i)\sigma_{LTD}^i}{\mu_i} \right) \geq MFR^i \dots \dots \dots (12)$$

Equation 12 is the first constraint for the second model. The model will have N different constraints like the above constraint for N different products.

The second constraint in the above model is where the aggregate type 2 service level or weighted average fill rate of the system is greater than or equal to AFR . The aggregate fill rate is $fr_i(z_i)$ multiplied by the ratio of the average demand μ_i of product i , to the total average demand of all products μ_{total} .

$$\text{Minimize cost: } \sum_{i=1}^N c_i z_i \sigma_{LTD}$$

Subject to

$$fr_i(z_i) \geq MFR^i \quad \text{for } i = 1 \dots N$$

$$\sum_{i=1}^N \frac{\mu_i}{\mu_{total}} fr_i(z_i) \geq FR$$

The objective function of this model is the same as the one for the first model. We are minimizing the inventory cost of the safety stock. We do not account for double counting of backorders in our model for fill rate as mentioned by Silver (1970) and others.

However Larsen et al (2007) have proved that for one to one continuous replenishment ($S-I, S$) policy, type 1 and type 2 service level are the same. Type 1 service level measures the number of stock-outs over a given planning period. It provides the

probability with which the new order received will be used to satisfy the demand over lead time with no stock-outs. In case of $(S-1, S)$ policy the new single unit received will be used to either satisfy the new demand or the backorder of the previous time period. Here measuring the fill rate would mean measuring the probability with which the incoming unit is used to satisfy the new demand, which is measuring type 1 service level again. In the implementation process we also verify the above observation computationally.

CHAPTER 6

SOLUTION TECHNIQUES

The models are solved by making use of the Lagrangian multiplier. The Lagrangian multiplier λ is introduced in the second constraint which links all the past decisions together. In the Lagrangian solution the equation is multiplied by $\frac{1}{C}$, where C is the total inventory holding cost of all the components. This is done to make sure that the solution obtained after carrying out the Lagrangian multiplier calculations is between 0 and 1. Only by obtaining a value between 0 and 1 we can get closed form solutions for z_i values. Note that due to the presence of natural logarithm (ln) in the solution to minimize the total cost it is divided by C , which is a constant.

After adding the Lagrangian multiplier and carrying out the first order differentiation we get, the following closed form solutions.

$$z_i = \text{Max} \left[\sqrt{-2 \ln \left[\frac{(\sqrt{2\pi}) [c_i \sigma_{LTD}^i \mu_{Total}]}{C \lambda \mu_i} \right]}, \Phi^{-1}(MSL^i) \right] \dots \dots \dots (13)$$

The above is for first model, for the second model we have

$$\Phi(z_i) = \text{Max} \left[- \left[\frac{c_i \mu_{Total}}{C \lambda} \right] + 1, fr_i^{-1}(MFR^i) \right] \dots \dots \dots (14)$$

Please refer to Appendix A and Appendix B for the complete calculation of the Lagrangian solutions for the two models respectively.

By taking the maximum value of z_i in the both the solutions we ensure that the minimum type 1 service level and fill rate are achieved respectively. The standardized z_i score is

found out by taking the inverse of the standard normal cumulative distribution function for $\frac{c_i \mu_{Total}}{c\lambda}$. Equation 14 is now modified as;

$$z_i = \text{Max} \left(\Phi^{-1} \left(- \left[\frac{c_i \mu_{Total}}{C\lambda} \right] + 1 \right), fr_i^{-1}(MFR^i) \right) \dots \dots \dots (15)$$

The models will be solved using Microsoft excel. There are two variables associated with the solutions of the two models, which are z_i and λ . The solution methodology starts with assuming $\lambda=0$. The corresponding z_i values are then obtained, and then are tested to see if they satisfy the constraints for the first and the second model respectively. Till the z_i values satisfy the constraints an iterative search of the Lagrangian multiplier λ , is being carried out. After finding out the z_i values the corresponding safety stock is calculated. In our present work the focus is to find out the appropriate safety stock under service level restrictions and the minimum inventory cost. The above closed form solution does not account for variability in the demand of the products quantified by σ . We anticipate therefore second model to perform poorly as compared to the first model.

The solution is a closed form solution for both the models. By observing the two solution equations the two service level measures will be impacted by volume of demand (μ_i), lead time (σ_{LTD}^i) and the cost of product (c_i). We anticipate service level measures to go up with increase in volume of demand and decrease with increase in lead time and cost.

CHAPTER 7

COMPUTATIONAL RESULTS

7.1 Data simulation

One of our primary goals is to achieve results that are applicable to real industrial set-ups.

We first start by developing a multi product inventory with characteristics that are similar to what is observed in industry. The data generation is carried out in Microsoft Excel.

The automated random number generation function in Excel is used for this purpose. We generate a two thousand product inventory with varying weekly demand distributions, lead times and cost. The multiproduct inventory ranges from products with slow demand (a mean demand of 0.1 per week) to fast demand (a mean demand of 100 per week).

Similar assumptions are also made for lead times and cost. The following table summarizes the important parameters associated with the multi product inventory developed.

Variable	Minimum Value	Maximum Value
Weekly demand	0.1	100
Lead time (weeks)	3	30
Cost (\$)	10	100000

Table 2 Data simulation variables

We also make the following important assumptions to ascertain that our simulated data represents a wide spectrum of products across different types of industries.

- 1) We ensure that there are enough products with slow moving demand, by dividing the two thousand product inventory evenly. We have a thousand products with mean weekly demand less than or equal 20 and remaining thousand with mean weekly demand greater than 20.
- 2) In industry the slow moving products typically have longer lead times. In case of fast moving demand the lead times are shorter. The slow moving products accordingly have a longer lead time while the fast moving products have a shorter lead time in the data simulated. The lead times go on decreasing from high to low as the demand increases from slow to fast.
- 3) The same is also true for cost. Slow moving products are likely to be costlier than the fast moving products and our data also reflects this generality.
- 4) The generalizations made in the previous two assumptions are applicable for the majority of the products in the market. There do exist a few products which are exception to the above rules. We also simulate data with products that have slow moving demand, which are cheaper and with shorter lead times to account for the above exception. Similarly it also makes sense to have fast moving products with longer lead times and high cost.

We now discuss the results to the numerical solutions obtained to the two models.

7.2 Model 1 Numerical Solution

The solution techniques used to obtain closed form solutions are for demand that follows a normal distribution. In industry a popular choice for approximating demand in forecasting is assuming that it follows a Poisson distribution. In our computational

results we make use of normal approximation of Poisson distribution. We test the robustness of our model by using this approximation and also make certain that the results are applicable to majority of products in the industry. The standard method to carry out this approximation can be found in general inventory management texts like Nahmias (2005), Silver & Peterson (1979). The value of mean is same in both normal and Poisson distributions. The standard deviation σ in the approximation is the square root of the mean. In our model the standard deviation calculation is as follows,

$\mu_i = \text{mean weekly demand for product } i.$

$$\sigma_i = \text{standard deviation of weekly demand} = \sqrt{\mu_i} \dots \dots \dots (16)$$

The solution starts with generating the mean demand (per week), lead times (in weeks) and the cost for the 2000 multi product inventory. The data is simulated based on the assumptions described in section 7.1.

- 1) After generating the demand we calculate the standard deviation for demand of product using equation 16 above.
- 2) The lead time in the solution is assumed to have a normal distribution with a specified mean and standard deviation.
- 3) The desired minimum type 1 service level for every product in multi product inventory is assumed to be 85 percent, while the weighted aggregate type 1 service level for the entire inventory is assumed to be 95 percent.
- 4) The average demand over the lead time is calculated using equation 1 while the standard deviation of demand over the lead time is calculated using equation 3: See section 4 for details.

- 5) The z_i values required to achieve desired type 1 service level is calculated using the Lagrangian solution mentioned in section 6, equation 14.
- 6) After obtaining the z_i value the type 1 service level for each product is found by calculating the corresponding cumulative distribution function $\Phi(z_i)$. This is done by using the NORMSDIST function in excel.
- 7) The corresponding weighted type 1 service level for each product is found out using the following equation,

$$\text{Weighted Type 1 service level} = \frac{\mu_i}{\mu_{total}} \Phi(z_i) \dots \dots \dots (17)$$

The aggregate weighted type 1 service level is found out by summing up all results found by equation 17 for all products.

- 8) The required desired safety stock is calculated by using standard equation 2 presented in section 4.
- 9) We are also interested in comparing the results obtained for the two models. The Fill rate is then calculated using the standard loss function equation.

$$fr_i(z_i) = \left(1 - \frac{\phi(z_i) - z_i(1 - \Phi(z_i))\sigma_{LTD}^i}{\mu_i} \right) \dots \dots \dots (18)$$

- 10) The final results were analyzed by classifying type 1 service levels and fill rates in different groups, for example number of products with fill rate from 75 to 80 and number of products with fill rate from 80 to 85 and so on.
- 11) An iterative line search for λ is carried out till an aggregate weighted type 1 service level of 95 percent is achieved. As mentioned above the aggregate weighted type 1 service level for 2000 products is found out by summing by the individual weighted type 1 service levels found by equation 17. We start with a

Lagrangian multiplier of zero and keep increasing it in either direction till desired results are achieved.

12) The fill rate calculation is also done by using equations used in the Sobel (2004) paper. We want to compare our fill rate equations against those that are similar to us.

The summary of results is given in Table 3 for the 2000 multi product inventory. For the 2000 product inventory an aggregate type 1 service level of 95 percent for the entire inventory is achieved at a Lagrangian multiplier value of approximately 32.

Model 1 results (Type 1 service level)	
Total demand μ_{total}	86288.89
Total cost C	\$60,012,696.58
Lambda λ	31.7
Total optimized cost	\$585,372,332.64
Total no of products in Safety stock	51946.9045
Desired Aggregate type1 Service level(ASL)	0.95
Aggregate type 1 service level obtained (ASL)	0.9498
Desired type 1 service level for individual products	0.95
Number of products with desired type 1 service level	657
Desired minimum type 1 service level for individual products	0.85
Number of products with desired minimum type 1 service level.	1377

Table 3 Model 1 results summary

The following figure is a classification of number of products in different classes of type 1 service level.




Figure 1: The number of products in different classes of Type 1 service level
 We inferred the following observations from the results obtained,

- The model is constrained by a minimum service level constraint. For the Lagrangian solution though all the products achieve a minimum service level, only one third of the 2000 products have the desired service level.
- The products with slow moving demand are the products with minimum service level while the products with fast moving demand are the ones with desired service level. The above two observations reflect the general theory that obtaining high service level's when demand is slow is difficult. The inventory is divided into half for both slow and fast moving products and all the slow moving products have a minimum service level. The service level of the products increase's as the nature of the mean demand per week changes from slow to fast. This increase can be also attributed to the second constraint in our optimization model. In the second constraint we are trying to achieve a weighted aggregate service level for all products. The ratio of mean demand for a product to the total demand is taken as the weight, and products with fast moving demand will clearly have higher weights. These results are illustrated in figure 2 below.




Figure 2: Change in Type 1 service level values with change in mean demand of the products.

- The change in cost also provides similar results. The objective function is trying to minimize the inventory cost while the constraints are trying to maximize the service level. The objective function and the constraint for service level are in conflict with each other. This reflects in minimum service level for costlier products and desired service level (or higher) for cheaper products as illustrated by the Figure 3 below.




Figure 3: Change in Type 1 service level values with change in cost of the products.

- The service level of the products is increasing with a decrease in lead time. We have assumed that products with longer lead times have slow moving demand and are costly. It needs to be examined whether this change in service level is a direct consequence of the change in lead time. We therefore look at products which are exception to the assumptions made i.e. fast moving products with longer lead times and costly. We observe that service level drops down only when increase in lead time is accompanied by increase in price. The number of these exceptional products is still small in comparison to the products that follow the general assumption. The above observation may be still impacted by the inventory data that we have generated. As discussed in Section 6, for a closed form and under equal circumstances for the three parameters lead time will have an indirect impact on type 1 service level. We confirm these observations by rigorously testing the model for changes in lead times under different cost circumstances.
- As expected increasing the value of the Lagrangian multiplier increases the service levels, though this comes off at the trade off of an increase in cost. By changing the multiplier the manufacturers can decide how to balance the inventory cost and desired aggregate service level they wish to obtain.
- In a $(S-I,S)$ inventory policy the type 1 service level and fill rate values should be equal to one another. The fill rate calculations by using the standard equations mentioned in section 4 are consistent with above proved theoretical result. The fill rate calculations by using the equations provided in the Sobel (2004) paper are not close to the type 1 service level values. In simulation we verify if the fill rate results obtained by using equations in section 4 are also accurate and consistent.

We observe that for the 2000 product inventory 1923 have fill rates of 75 or more will the remaining have fill rates less than 75. The following figure gives exact number of products in different ranges of fill rates.

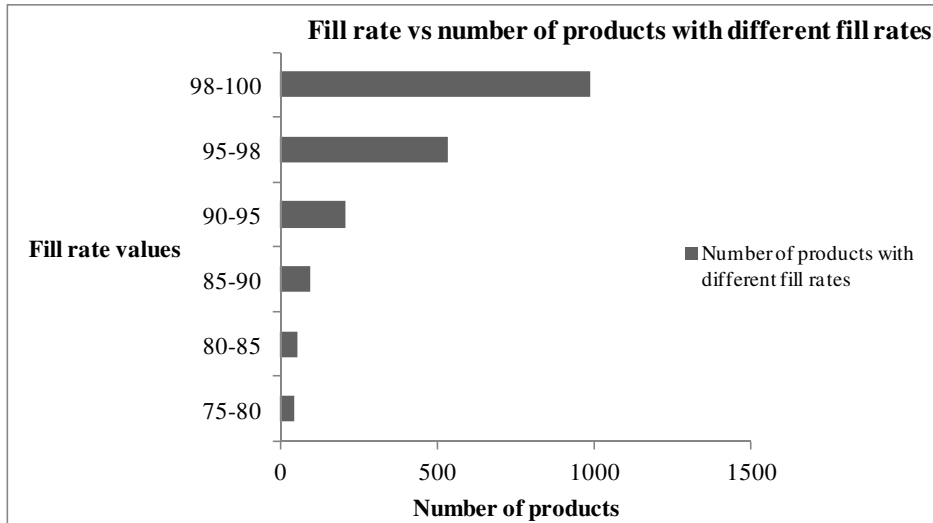


Figure 4: The number of products in different classes of Fill rates

- The type 1 service level and fill rates values do not match for slow moving products. For fast moving products they match perfectly. In general type 1 service level values are higher than that of fill rates. This difference is evident in extremely slow moving products. We can account this difference to double counting of backorders. Double counting of backorders is more evident in products that have very slow moving demand and is nonexistent for products with high moving demand. In the results we observe that as the nature of the demand starts to change from slow to fast the fill rates also start to increase. We also observe as the nature of demand changes the type 1 service level and fill rates start to match. In simulation we account for double counting by measuring only the inventory left after satisfying backorders towards fill rate calculations. This is

described in more detail in the preceding sections. The fill rates for fast moving products match with the type 1 service level though it needs to be rigorously tested by simulating different inventory scenarios for different types of products.

In section 8 we simulate different types of demands to further analyze the impact of these variables on type 1 service level.

7.3 Model 2 Numerical Solution

The numerical results for Model 2 are also obtained using similar steps as that for Model 1. The only major difference is the calculation of standardized z_i score for each product. The z_i score is calculated using the Lagrangian solution obtained by equation 15 mentioned in section 6. The calculations for weighted aggregate fill rate are also computed by summing all the individual weighted fill rates for all the products. The same weight of $\frac{\mu_i}{\mu_{total}}$ is used for the Model 2. Subsequently the new minimum fill rate is 85 percent while 95 is the desired fill rate value.

The following table summarizes the results obtained for Model 2.

Model 2 results (Fill rate)	
Total demand μ_{total}	86288.89
Total cost C	\$60,012,696.58
Lambda λ	175
Total optimized cost	\$509,101,286.64
Total no of products in Safety stock	51723.19083
Desired aggregate fill rate for all products (<i>AFR</i>)	0.95
Aggregate fill rate obtained for all products (<i>AFR_{obtained}</i>).	0.9501
Number of products with desired fill rate (<i>FR</i>)	2030
Desired aggregate type1 service level(<i>ASL</i>)	0.95
Aggregate type 1 service level obtained (<i>ASL_{obtained}</i>)	0.89

Number of products with desired Type 1 service level	1135
--	------

Table 4: Model 2 results summary

The following are the observations from the results obtained for the second model,

- The safety stock values and service level results obtained for Model 2 are very similar to those obtained by Model 1. Larsen et al (2007) in their paper on comparing fill rate definitions for a base stock inventory control system prove in their theorem that for an inventory policy with base stock level S the Order fill rate or type 1 service level is the same as the Volume fill rate or type 2 Service level. The results obtained from the solution are consistent with the above theorem. As mentioned in the table above there are 1600 products with a type 1 service level greater than 75 and while all the 2000 products have a fill rate greater than 75.
- The type 1 service levels values and fill rates though a close match to one another, the aggregate type 1 service level of all the products is not close to 95 percent. This confirms our deduction that Model 2 provides poor results as compared to the first model.
- The fill rate values obtained from the solution of Model 2 are slightly different than those obtained from the solution of Model 1. This difference is only for slow moving products. For fast moving products, the fill rates obtained from both models match closely. The increase in the fill rates observed in solution of Model 2 can be attributed to increase in the amount of safety stock of slow moving products. Model 2 is constrained by minimum fill rate constraint forcing the safety stock values of slow moving products to go up.

- The total optimized cost obtained in Model 2 is lower than the total optimized cost of Model 1. The above difference is on account of obtaining lower safety stock values for fast moving products in Model 2.

To summarize our results for the two Models, cost and nature of the demand seem to have a major impact on service levels. Though lead time individually does not have a drastic impact on service levels we need to verify the above deductions with further thorough computations.

The demand values in the solution are mean or average values of demand. The mean demand from week to week or for different time periods is constantly fluctuating. These inventory levels changes certainly impact the service levels across a wide planning horizon. The service level values obtained from the solution have to stay constant across a planning horizon. It is necessary to test the safety stock values for changing inventory levels, spread across a wide time horizon.

The fill rate values for slow moving products in Model 1 are less than that of type 1 service level. In an actual industrial setting this may have a significant impact on inventory performance as having reasonable fill rate is equally important. In a real industrial setting and with no double counting of backorders we expect a minor increase in fill rate values of the slow moving demand. This also needs to be verified by simulating different inventory scenarios. Hence the safety stock values have to be examined to account for all these uncertainties that a standard inventory in practice is susceptible too. The results obtained by the two models are similar. The first model

provides more accurate results. We thus focus our computational simulations for Model 1 only. For Model 2 even though we do not meticulously analyze complete details we examine the results obtained by our solution under a real inventory setting. In the succeeding sections we describe the simulation assumptions and report the simulation results.

CHAPTER 8

SIMULATION RESULTS

In this section we first start with experimental setup of the simulation and underlying assumptions made for it. We then provide the results associated with the simulation of nine basic products. We summarize the results of the simulation subsequently.

8.1 Simulation setup, assumptions and steps

We carry out our simulation in Microsoft Excel. Excel is readily available and well known data processing tool. It is easy to understand and also has Visual Basic as an inbuilt programming language. The presence of Visual Basic offers flexibility of coding different types of demand distributions and standard mathematical equations along with directly using standardized functions. It is a common tool used not just in Inventory control but also across other branches of Operations Research and business in general.

We make a few computational changes from the numerical solution in our simulation setup to ascertain that the safety stock obtained by the solution is applicable in an actual industrial setting. Along with the changes we also list the assumptions, simulation setup and steps followed in it.

- 1) In the simulation time is discretized in units corresponding to weeks. We thus consider weekly demand. The demand generated is for every week, the lead time values are in weeks and the type 1 service level is calculated per week.

- 2) The simulation is carried out for each product individually and the demand is simulated assuming it follows a Poisson distribution.
- 3) In our original solution we have assumed that lead times also have a normal distribution. This uncertainty in lead time logically reflects an increase of safety stock. This added complexity may not be helpful in testing the robustness of the safety stock values. Hence we relax the assumption of normally distributed lead time and assume a fixed lead time.
- 4) We have assumed a $(S-I, S)$ inventory policy. S here is the base stock policy value calculated by multiplying mean demand with the lead time for product i and then adding the safety stock value obtained from the solution in the previous section.

$$S = (\mu_i * L_i) + (Safety\ stock)_i \dots \dots \dots (19)$$

The calculation is represented by equation 19 mentioned above. The above calculations are carried in terms of week. The mean demand is assumed to be mean demand per week with the lead time also in weeks.

- 5) The nature of $(S-I, S)$ policy is one to one replenishment. It means the moment inventory level goes down below the base stock level S , we order to bring it up to back to level S . It simply means every time a demand is realized you place an order as you are always facing customer demand. Every time the inventory drops down with the number of units demanded for a given week. To bring the inventory level back up in next week we therefore order the exact number of units demanded.
- 6) The inventory level changes take place at the beginning and the end of the week. At the beginning of each week i , a demand is realized and the existing inventory

is used to satisfy this demand. The inventory level drops down and hence an order of same number of units is placed to bring the inventory level back up. This order has a lead time of L weeks. At the end of the week the order placed at the end of $i-L$ week arrives and inventory level goes up again. The inventory level is therefore updated at the end of each week. In literature inventory changes take place at the beginning of a time period t . The orders are placed at the beginning of a period, when the corresponding demand of that time period is still unknown. The order is also received after $t+L$ periods and also at the beginning of the period $t+L$. We do not use the general assumption in literature to make our work consistent with actual industrial practices.

- 7) The existing inventory is first used to satisfy a previous backorders for a given week $i-1$ to account for backorders only once. After accounting for backorders the leftover inventory is used to satisfy the demand for the week i . If there are no units left after satisfying the backorders of the previous week i , the type 1 service level is zero. Type 1 service level is also measured as zero if with the leftover inventory there are not enough units to satisfy the entire demand of week i . If the entire demand for week i is satisfied with existing inventory the type 1 service level is measured as one. Similar calculations for type 1 service level is done for all weeks in the entire planning horizon. The type 1 service level for the entire planning horizon is then calculated by dividing the number of weeks with no backorders over the total number of weeks in the planning horizon. The following equation represents the described calculation,

$$\text{Service level} = \frac{\text{Number of weeks with no backorders (Service level one)}}{\text{Total number of weeks in a planning horizon}} \dots \dots \dots (20)$$

8) To calculate fill rate we similarly first account for backorders of the previous week $i-1$. Then with leftover inventory entire or partial demand for week i is satisfied. In this way demand satisfied for each week is noted. If the entire inventory is used to satisfy demand of previous week $i-1$, then demand satisfied for week i is zero. If there is inventory left after satisfying backorders of previous week $i-1$, then demand satisfied for week i is either the full demand or a fraction of it. The fill rate is calculated by summing this demand satisfied across all weeks in the planning horizon and then dividing this summation over demand for the entire planning horizon. This calculation is somewhat similar to Silver (1970) equation as they take the ratio of demand satisfied over a fixed order quantity R . We formalize the steps described above in the following expressions. We first define the following variables. Let

$$\text{Demand satisfied for week } i = s_i$$

$$\text{Inventory left for week } i \text{ (after satisfying backorders for week } i - 1 \text{)} = I_i$$

$$\text{Demand observed for week } i = D_i$$

$$s_i = \left\{ \begin{array}{ll} 0 & \text{if } I_i = 0 \\ I_i & \text{if } D_i > I_i > 0 \\ D_i & \text{if } I_i > D_i > 0 \end{array} \right\} \dots (21)$$

$$\text{fill rate for a planning horizon of } n \text{ weeks} = \frac{\sum_{i=1}^n s_i}{\sum_{i=1}^n D_i} \dots (22)$$

9) We also calculate the average weekly fill rate along with calculating the fill rate for the entire planning horizon. The weekly fill rate is simply the ratio of demand satisfied for a week i which is s_i to the demand observed for the week i or d_i . We take the average of weekly fill rate for the entire planning horizon.

- 10) The input variables to the simulation experiment are the mean demand per week, the lead time and the safety stock for that particular product under consideration obtained from the model solution described in the previous section.
- 11) We start with the simulation for an infinite planning horizon. The infinite planning horizon consists of 10,000 weeks. In actual practice there takes a certain amount of time for inventory to start operating under steady state conditions. We start our simulation with inventory on-hand equal to base-stock level S , that is the entire inventory in the system under steady conditions. In the first few weeks thus the inventory is always high and type 1 service levels and the fill rates will be a hundred percent. We provide a warm up period in our simulation to account for these initial inventory dynamics. Type 1 service level and fill rate calculations start after this warm up period. We consider the time span of first 350 weeks as warm up period. The simulation is run for a minimum of twenty trials as specified in Law and Kelton (2000). Similar observations are found in the work of Kwon et al (2006) and Fu et al (1992).
- 12) The simulation starts with specifying the number of trials. For each trial demand for different weeks is generated, the corresponding inventory level changes are updated and calculations for type 1 service level and fill rate are carried out. This is carried out for all the weeks in a given planning horizon.
- 13) The results include the type 1 service level, fill rate, the mean demand and the standard deviation for each trial.
- 14) The average type 1 service level and fill rate for twenty trials is computed along with the distribution of error in the results obtained for these values.

8.2 Simulation results for 9 basic products

We start our simulation experiments for 9 products of different demand, lead time and costs chosen from the 2000 product inventory system optimized using Model 1. The primary goal of our simulation is testing the safety stock values obtained through Model 1 in a real inventory setting and trying to understand the impact of other factors on it. The following table summarizes the important variables of these 9 basic products,

Product No	Variables		
	Mean demand per week (number of units)	Lead time (weeks)	Cost (dollars)
1	1	17	\$44,424
2	12	9	\$29,595
3	29	7	\$37,233
4	47	6	\$7,637
5	68	3	\$15,496
6	87	3	\$438
7	99	4	\$456
8	99	1	\$356
9	0.1	27	\$52,713

Table 5 Mean demand, lead times and cost of the 9 products for simulation

The 9 products, though randomly chosen from the 2000 product inventory follow a generic trend of increase in the mean demand with decrease in lead times and cost. We

also simulate an extremely slow moving product with mean demand of 0.1 per week. In the following Table 6 we provide the first ten weeks simulated for product 3. The simulation table is for weeks after the warm up period. The backorders for the week of 344 are 13, which makes the total demand for week 344 as 42 and with an existing inventory of 38 generating 4 backorders. Similarly these backorders are added to next week demand.

Simulation Table									
Cycle/ Week	Demand/Order for week	Lead time weeks	Starting inventory of the day	Inventory after satisfying Backorders	Ending inventory of the day	Inventory left after satisfying Backorders	Demand satisfied	Type 1 service level	Type 2 Service level or Fill rate
344	29	7	38	0	30	24	24	0	0.85
345	28	7	30	0	28	25	25	0	0.90
346	18	7	28	7	36	25	18	1	1
347	25	7	36	11	42	36	25	1	1
348	32	7	42	10	53	42	32	1	1
349	32	7	53	21	54	53	32	1	1
350	30	7	54	24	53	54	30	1	1
351	17	7	53	36	64	53	17	1	1
352	27	7	64	37	55	64	27	1	1
353	23	7	55	32	57	55	23	1	1

Table 6: Example of Simulation carried out in Excel. Example is for Product 3 and for the first 10 weeks out of 10000 week planning horizon.

In our simulation results we expect to obtain type 1 service levels and fill rate values close to the values predicted by the numerical solution of our Model 1. In Table 7 we have provided the comparison of type 1 service levels and fill rates obtained by the simulation to the ones expected by the model solutions.

Product No	Type 1 service level from solution.	Type 1 service level from simulation.	Fill rate from solution.	Fill rate from simulation. For entire horizon.	Fill rate from simulation. Weekly Fill rate.
1	0.85	0.86	0.93	0.92	0.92
2	0.85	0.87	0.96	0.95	0.94
3	0.85	0.86	0.97	0.96	0.96
4	0.94	0.93	0.99	0.98	0.99
5	0.91	0.93	0.98	0.99	0.99
6	0.99	0.99	0.99	0.99	0.99
7	0.99	0.99	0.99	0.99	0.99
8	0.99	0.99	0.99	0.99	0.99
9	0.86	0.86	0.79	0.74	0.74

Table 7 Comparison of Type 1 service level and Fill rates obtained by simulation to expected solution.

The total results are also summarized in the figure 5 and in figure 6 for type 1 service level and fill rate respectively. The simulation results for each product represent the average type 1 service level and fill rates across the entire planning horizon for over twenty replications of 10000 weeks.

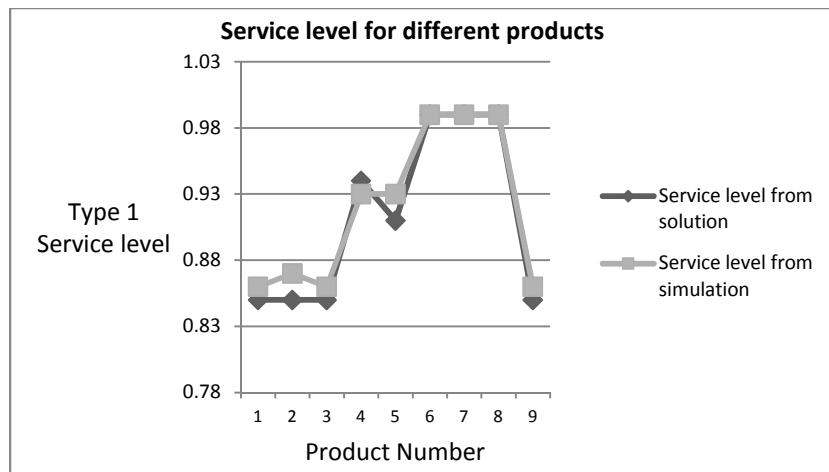


Figure 5: Service level comparison for nine basic products.

The average type 1 service level for 9 products is close to the value estimated in the computational solution.

For slow moving products we observe a higher service level than that predicted by the solution in Model 1. As the nature of the demand moves from slow to fast the service levels obtained from the simulation and solution are a perfect match. Product 9 is an extremely slow moving product and the service level drops down to the minimum required value; despite this the results obtained by the simulation and solution are almost a perfect match. We now look at the fill rate comparisons obtained for the nine different products using simulation. The average weekly fill rates are almost identical (they differ for product 2) to the fill rates for the entire planning horizon. The comparison with model solution is summarized in figure 6 below. The fill rates values used here are the fill rates for the entire planning horizon. The fill rate values obtained in the simulation are close to the values evaluated in the solution.

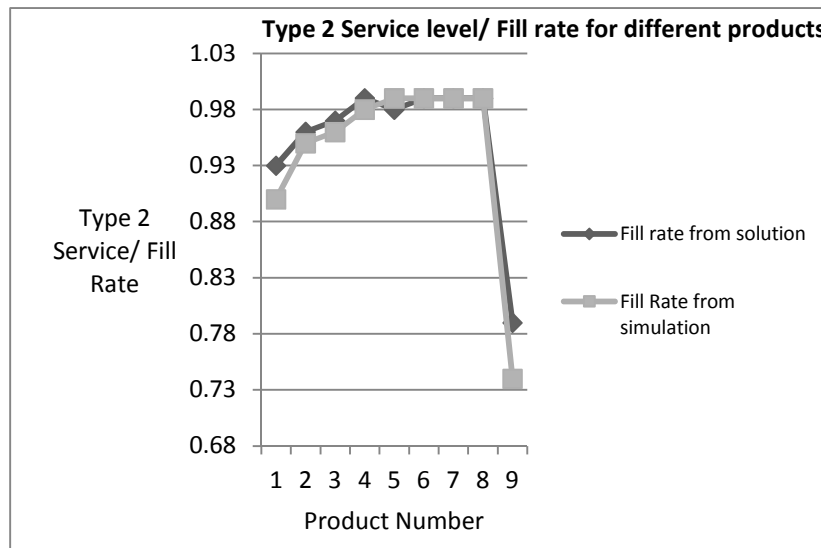


Figure 6: Fill rate comparisons for nine basic products.

The fill rates obtained by the solution and simulation start to perfectly match as the nature of the demand moves from slow to fast. For slow moving products there is a minor difference in the fill rates with the ones obtained from solution slightly higher than that of the simulation. For the extremely slow moving product 9 the fill rate drops down as guessed though the simulated value is very close to the numerically computed value. To validate our assumptions about cost and lead time we arrange the products in order of increasing lead times and cost, plot the two service levels as represented in figure 7 and figure 8 respectively. We also gain an understanding of the following important insights.

- The simulation results confirm our deductions that changes in cost impact service levels. We had assumed that slow moving products in our multi product inventory generally have higher lead times and are costlier. It also becomes necessary to thoroughly examine the products which are exceptions to this assumption. We need to examine the model to see the kind of service levels obtained for slow moving products which are cheaper and for fast moving products which are costlier.




Figure 7 Change in Type 1 service level and Fill rate with change in cost.

- The changes in lead times do not have an individual impact on service levels. From the results we see that increase in lead times seems to negatively impact service levels. However the 9 products chosen for simulation also have a corresponding change in cost and mean demand with changes in lead time. The lead time decreases with increasing mean demand and increases with the cost. This confirms our deduction that changes in lead time do not have an individual impact on service level changes. We further confirm this observation in our simulation. There are also products in the market which are exception to assumptions we have made for lead times and cost. There are products with fast moving demand and lead times which are not necessary very short. We further need to test the applicability of our results by carrying out simulations for products with fast demand and longer lead times and for slow moving products with shorter lead times.




Figure 8 Change in Type 1 service level and Fill rate with change in lead times

- The average service level values obtained by the simulation are very close to the expected solution values. The average fill rates obtained by simulation are

marginally less than the expected values. This may be on account of rounding errors as lead times and safety stocks results are in fraction. There is however noteworthy variability for service levels and fill rates obtained from trial to trial. This variability is more noticeable in the case of slow moving products. The random number generator in Excel may be responsible for this variation. The histogram in figure 9 and subsequent table provide the variation in the results obtained from the twenty trials. It would be interesting to see how simulation results change if the demand is generated using another statistical tool.

- To calculate the above distribution of error for service levels obtained from simulation we calculate the standard deviation of the twenty replications. These standard deviation values are provided in the table below for 9 products.

Product No	Variables				
	Mean demand	Type 1 service level simulation	Fill rate simulation	Standard deviation for Type 1 service level across 20 replications	Standard deviation for Fill rate across 20 replications
1	1	0.86	0.92	0.0133	0.0150
2	12	0.87	0.95	0.0087	0.0058
3	29	0.86	0.96	0.0060	0.0025
4	47	0.93	0.98	0.0043	0.0009
5	68	0.93	0.99	0.0035	0.0003
6	87	0.99	0.99	0.0002	1.92×10^{-5}
7	99	0.99	0.99	0.0003	2.6×10^{-5}

8	99	0.99	0.99	0.0004	2.6×10^{-5}
9	0.1	0.86	0.74	0.0157	0.0250

Table 8 Standard deviations in Type 1 service level values and Fill rates for 20 replications of 9 products.

As seen in the table above the error in the simulation results is very small for products with high demand. We also had assumed that products with high demand are cheaper and have low lead time. Hence we can assume products with low demand and low cost have less error in the simulation service level measures. We though verify these deductions in the next section.




Figure 9: Variability in results for the 20 simulation replications of Product 1 in Excel.

The following table lists the important statistical variables indicating the variation across each replication for all the nine products.

Product No	Mean	Median	Standard Deviation	Skewness	Minimum	Maximum	1 st Quartile	3 rd Quartile
1	0.8690	0.8671	0.0107	0.28	0.8494	0.8893	0.8628	0.8893
2	0.8717	0.8703	0.0087	0.13	0.8540	0.8905	0.8654	0.8905
3	0.8606	0.8603	0.0060	-0.60	0.8470	0.8709	0.8573	0.8709
4	0.9342	0.9339	0.0043	-0.39	0.9244	0.9409	0.9306	0.9409
5	0.9313	0.9319	0.0035	-0.10	0.9250	0.9391	0.9293	0.9319
6	0.9995	0.9995	0.0002	-0.45	0.9990	0.9990	0.9995	0.9999
7	0.9988	0.9987	0.0003	-0.04	0.9981	0.9993	0.9985	0.9993
8	0.9987	0.9988	0.0004	-0.29	0.9980	0.9994	0.9983	0.9994
9	0.8646	0.8663	0.0142	-0.26	0.8380	0.8907	0.8558	0.8907

Table 9: Mean, Median, Skewness, Quartiles and Range of Type 1 service level values for 9 products across 20 replications.

- We categorized demand of less than a mean demand of 20 products as slow moving demand. In the results only Product 1 and Product 9 with mean weekly demand of 1 and 0.1 respectively have characteristics of slow moving demand. There is also a significant range in the lead times generated and cost. This huge range does not justify classification of the products only as high or low. To clearly understand the impact of the above variables on service levels particularly for slow moving products a much simpler and clear classification of these variables is required.

The observations lead to a more detailed classification of products and need to test safety stock of these products by simulation. In the next section we describe the simulation of twenty seven products to account for all combinations of demand, cost and lead time.

8.3 Simulation results for 27 classified products

We categorize products into low, medium and high with this categorization being done for nature of demand, lead time and cost. We test all the products that can be categorized by a combination of these variables. The simpler classification of products is in table 10, below.

	Variable		
	Mean demand per week	Lead time (weeks)	Cost (\$)
Low	0.1	3	10
Medium	10	10	100
High	100	30	1000

Table 10 Mean demand, lead times and cost for our experimental design.

The above classification of products characteristics leads to 27 products with different combinations of nature of demand, lead time and cost. The 27 different product combinations are described in table 11 below.

Product no	Variables	
	(Demand per week/Lead time/Cost)	
1	(0.1/3/10)	(Low/Low/Low)
2	(0.1/3/100)	(Low/Low/Medium)
3	(0.1/3/1000)	(Low/Low/High)
4	(0.1/10/10)	(Low/Medium/Low)
5	(0.1/10/100)	(Low/Medium/Medium)
6	(0.1/10/1000)	(Low/Medium/High)
7	(0.1/30/10)	(Low/High/Low)
8	(0.1/30/100)	(Low/High/Medium)
9	(0.1/30/1000)	(Low/High/High)
10	(10/3/10)	(Medium/Low/Low)
11	(10/3/100)	(Medium/Low/Medium)
12	(10/3/1000)	(Medium/Low/High)
13	(10/10/10)	(Medium/Medium/Low)
14	(10/10/100)	(Medium/Medium/Medium)
15	(10/10/1000)	(Medium/Medium/High)
16	(10/30/10)	(Medium/High/Low)

17	(10/30/100)	(Medium/High/Medium)
18	(10/30/1000)	(Medium/High/High)
19	(100/3/10)	(High/Low/Low)
20	(100/3/100)	(High/Low/Medium)
21	(100/3/1000)	(High/Low/High)
22	(100/10/10)	(High/Medium/Low)
23	(100/10/100)	(High/Medium/Medium)
24	(100/10/1000)	(High/Medium/High)
25	(100/30/10)	(High/High/Low)
26	(100/30/100)	(High/High/Medium)
27	(10/30/1000)	(High/High/High)

Table 11 Variable information of the 27 classified products.

For this simulation we use Minitab, a statistical software to generate mean weekly demand. We are inspecting if part of variation in our results from replication to replication is on account of the random number generator in Excel. We again observe that average weekly fill rates and fill rates for the entire planning horizon are very similar, identical for most products. The results are consistent with those from the previous section. We compare the type 1 service level and fill rates for the 27 products in table 12 below.

Product No	Type 1 service level from solution.	Type 1 service level from simulation.	Fill rate from solution.	Fill rate from simulation. Fill rate for entire horizon	Fill rate from simulation. Weekly Fill rate
1	0.99	1.00	0.99	0.97	0.97
2	0.99	0.96	0.97	0.98	0.98
3	0.85	0.75	0.39	0.69	0.69
4	0.99	0.98	0.98	0.98	0.98
5	0.98	0.98	0.97	0.94	0.93

6	0.85	0.92	0.66	0.77	0.76
7	0.99	0.99	0.99	0.99	0.98
8	0.96	0.97	0.94	0.93	0.93
9	0.85	0.82	0.80	0.78	0.76
10	0.99	1.00	0.99	1.00	1.00
11	0.99	1.00	0.99	1.00	0.99
12	0.85	0.84	0.93	0.95	0.96
13	0.99	1.00	0.99	1.00	1.00
14	0.99	1.00	0.99	1.00	1.00
15	0.85	0.85	0.96	0.93	0.93
16	0.99	1.00	0.99	1.00	1.00
17	0.99	1.00	0.99	1.00	1.00
18	0.85	0.84	0.98	0.90	0.91
19	0.99	1.00	0.99	1.00	1.00
20	0.99	1.00	0.99	1.00	1.00
21	0.96	0.96	0.99	1.00	1.00
22	0.99	1.00	0.99	1.00	0.99
23	0.99	1.00	0.99	1.00	0.99
24	0.92	0.96	0.99	0.99	0.99
25	0.99	1.00	0.99	1.00	0.99
26	0.99	1.00	0.99	1.00	1.00
27	0.85	0.85	0.99	0.96	0.95

Table 12 Comparison of Type 1 service levels and Fill rates for the 27 classified products. The results obtained by simulation compared with the expected model solutions.

The nature of demand has a direct relationship with service levels while the increase in cost inversely impacts service levels. The suggestion that change in lead time does not heavily impact service level is apparent in these results. The variability of type 1 service level and fill rate for different trials is less when using data from Minitab. We start our analysis by first comparing type 1 service level obtained by the simulation to the expected values from solution. The results are presented in figure 10 below.




Figure 10 Type 1 service level comparison for 27 products.

The simulation results match closely with the results obtained from the solution again. For the slow moving products there is a slight difference in the service level values though as the nature of demand moves from slow to fast the results of simulation again match perfectly. For products with medium demand that is mean weekly demand of 10 units higher service levels are obtained for cheap products. For all costly products as predicted by the solution only minimum service levels are obtained in the simulation. We observed similar results when the fill rates for the 27 products obtained by simulation were compared with the solution. In the following figure we are comparing the fill rates across the entire planning horizon for the 27 products.

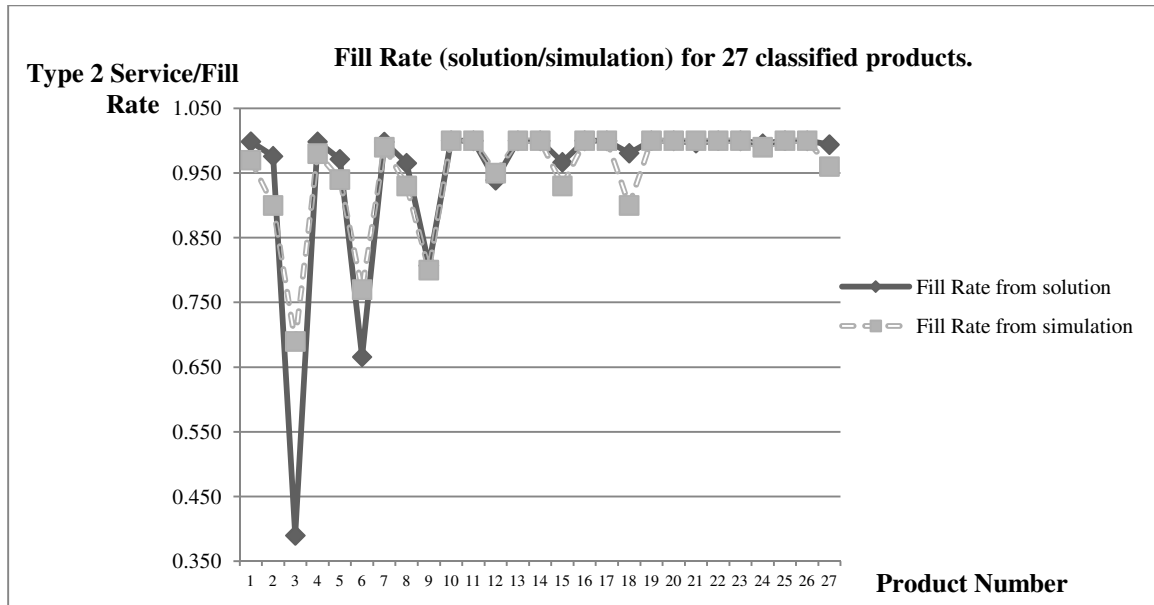


Figure 11 Fill rate comparisons for 27 products. Fill rates obtained by simulation compared with expected Fill rates solution.

Fill rates obtained by simulation for slow moving products are higher than the expected values from the fill rates. The fill rates obtained for products with medium to fast products are close to that expected from the solution. We also compare the type 1 service level and fill rates obtained by simulation. We observe that type 1 service levels and fill rates are similar. The results are presented in figure 36 in Appendix C. We first look at the change in service level with change in nature of demand. These corresponding results are provided by figure 12, 13 and 14. Consistent with our primary observations we confirm for slow moving and costly products, lower service levels are obtained. In figure 12 we see the impact of change in volume of demand and cost on service levels for products with low lead time. To clearly understand the particular impact of each variable we vary only two variables and keep the third constant. For the first case we change mean demand per week and cost while keeping lead times constant. We also note that for costly products lower service levels are obtained with all types of lead times. This

affirms our deduction that there is not a large impact on service level associated with changes in lead time.

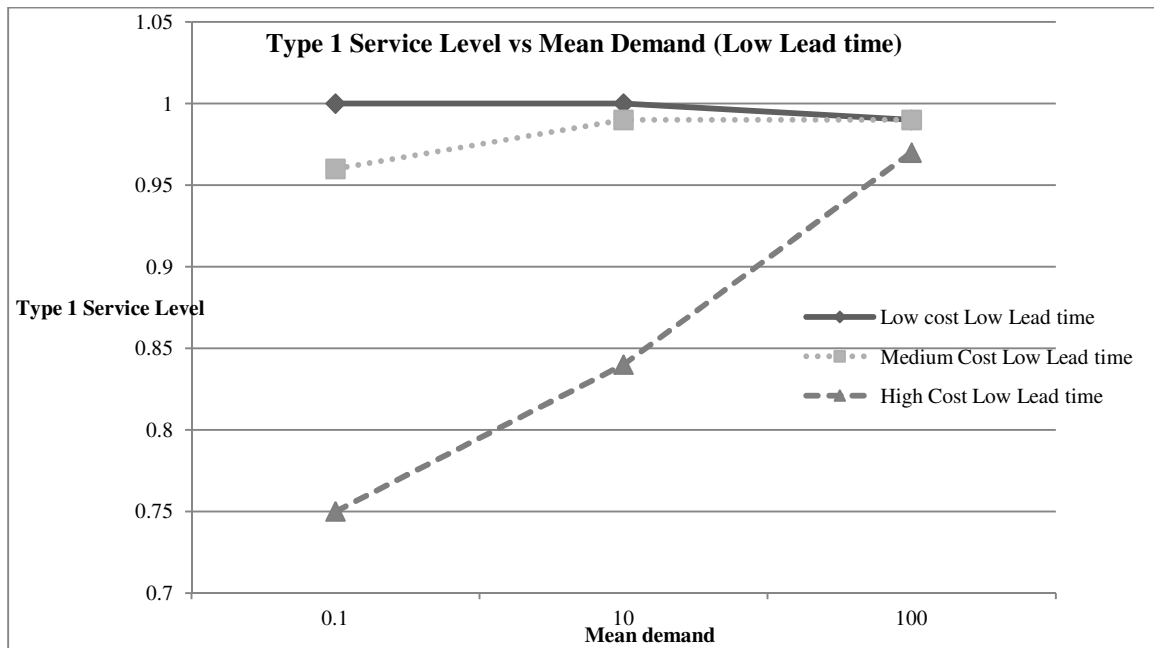


Figure 12 Change in Service level, with change in mean demand and change in cost. (Low lead times)

Similar observations are recorded for change in nature of demand and cost for products with medium and high lead times. The corresponding figures are attached in Appendix C.

We next look at the impact of change in lead time on type1 service level. We keep cost constant here and change nature of demand and lead time. We unite the three different graphs for a concise representation and for a simpler understanding of results.

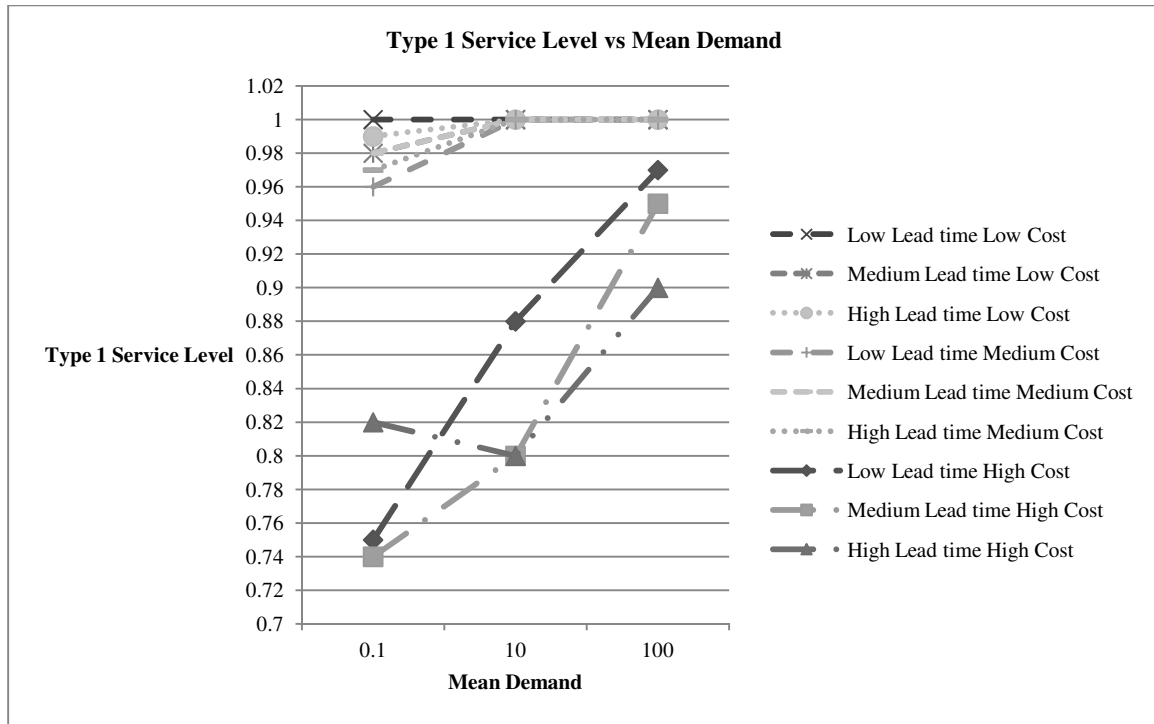


Figure 13: Change in Type 1 service level, with changes in demand for different lead times/cost product characteristics.

For the three categorizations of lead times (low, medium and high) low service levels are obtained only at high cost. These service level values go on increasing as the nature of demand moves from slow to fast. For products with low to medium lead times higher type 1 service levels are obtained for cheaper and medium cost products. This is consistent with standard inventory theory that lower lead time provides as a buffer to the uncertainty in the nature of demand. The individual graphs for different costs are attached in Appendix C. From Figure 13 we also conclude that change in cost impacts both fast moving and slow moving products. In figure 13 we see for Product with High lead time and High cost a decrease in type 1 service level values for medium demand. The type 1 service level for products with similar cost and lead time structure is slightly higher if the demand is low. This difference can be accounted for variation in mean demand which is represented by standard deviation σ in our work. As we are assuming normalized Poisson

demand the variation is calculated by taking square root of the mean demand. The low demand we have assumed is 0.1 and the variation of this demand ends up higher than the mean demand. For this increase in variation the model provides higher safety stock values. As similar conditions are not present when the mean demand increases we see a small decrease in service level values obtained. For slow moving products even with low lead times lower service levels are obtained for costly products. The same results can again be verified by purely changing lead times and cost by keeping the nature of demand constant. The results are represented in figure 14 below.

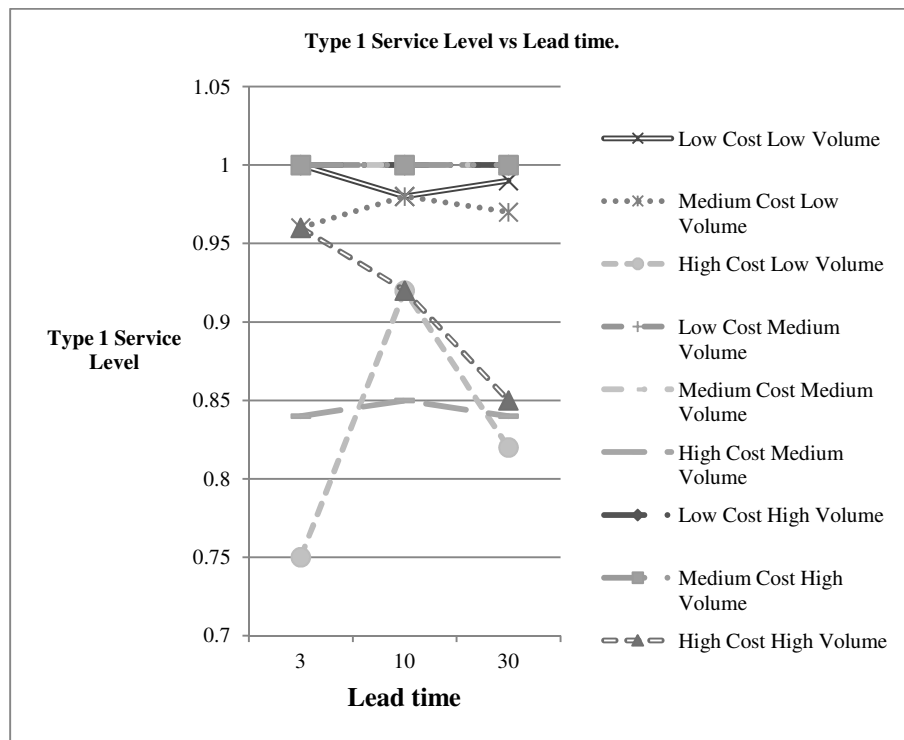


Figure 14 Change in Service level with change in lead time for different cost/demand product characteristics.

Similar observations can be noted if we vary cost for different types of demand and for given sets of lead times. We expect to see the service levels going down as cost is increasing. The corresponding results are provided in figure 15. All the lines have a

downward shape confirming an inverse relationship between cost and service level. The fill rate values obtained by the simulation are very close to the type 1 service level values. The change in the nature of the above three variables has almost identical impact on fill rates for the 27 products. The fill rates vary with change in nature of demand and cost and are independent of the lead times. In figure 16 we provide the unified graph for changing nature of demand and cost for given sets of lead times. The other corresponding results are provided in the Appendix C.

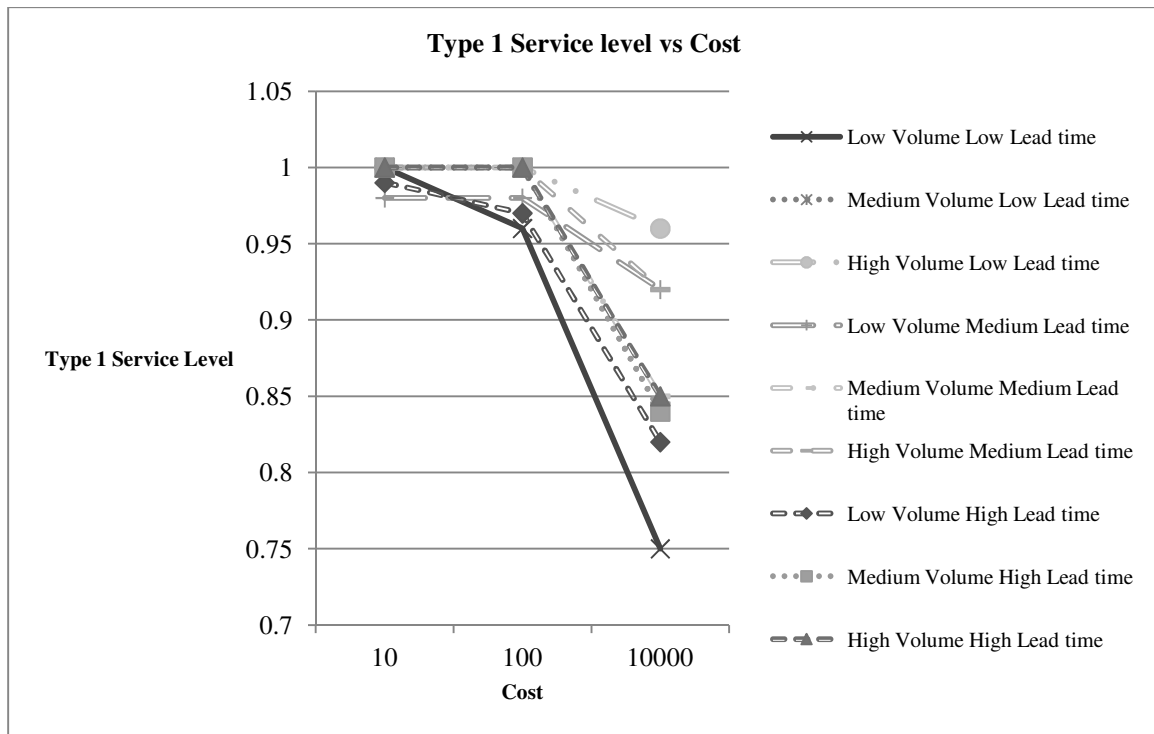


Figure 15: Change in Service level with change in cost for different volume/lead time product characteristics.

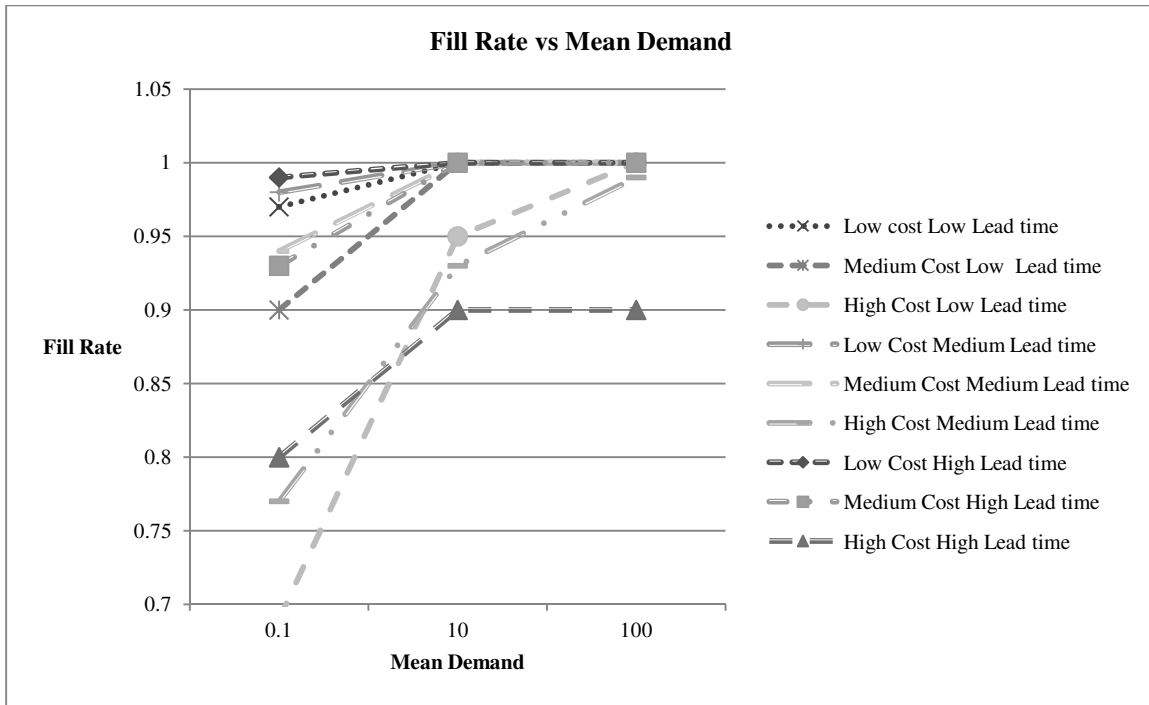


Figure 16 Change in Fill rate with change in the volume of demand and cost, for different lead times.

We now look at the error in simulation results with changes in the above variables. We use standard deviation as an estimate of error. We calculate the standard deviation in the type 1 service level values and plot it against corresponding changes in volume of demand, lead time and cost. We first start by observing graphs for standard deviation in type 1 service level values. The x axis has a non linear scale being divided into 3 parts. The x axis represents mean demand. In this case the first part is of the first nine points all of which have x coordinate as 0.1, the next nine have it as 10 and finally the last nine have it as 100. Similarly the other graphs for lead time and cost are set up.

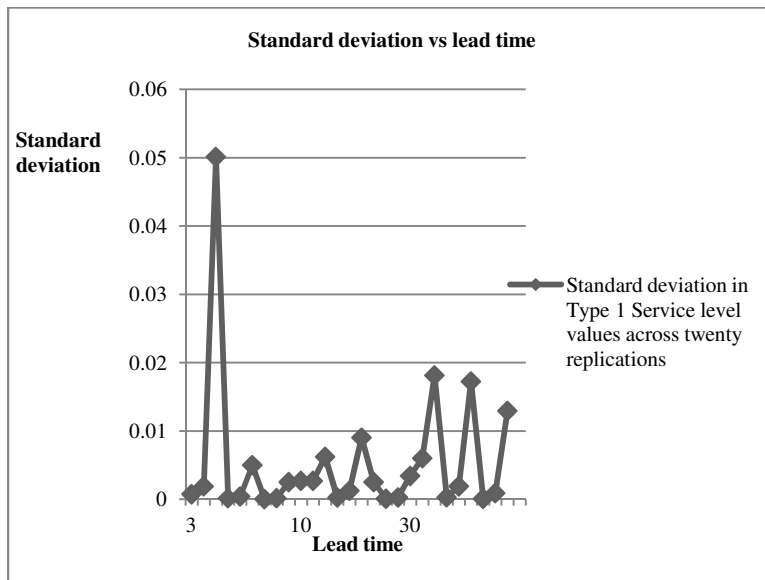



Figure 17 Changes in standard deviation of Type 1 service level values with changes in volume of demand and lead time respectively.

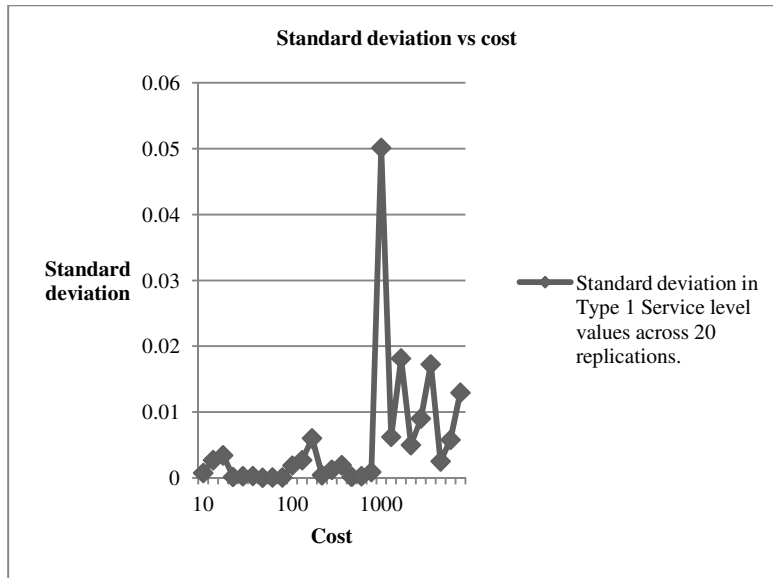


Figure 18 Changes in standard deviation of Type 1 service level values with changes in cost.

From the two figures above we see the error in simulation is higher in costly products and from those the highest deviation occurs for slow moving products.




Figure 19 Changes in standard deviation of Fill rate with changes in cost.

The standard deviation values for twenty replications of fill rates also change with change in cost. We provide the change in fill rate versus cost here will for other graphs in Appendix C.

To summarize the results of above simulations;

- In our model we have placed a higher weight on products with fast moving demand. This higher weight clearly reflects on obtaining better type 1 service levels and fill rates for fast moving products. For slow moving products only minimum service levels are obtained because of the objective function minimizing cost and on account of lesser weight being placed on them.
- The models are driven by the nature of the demand and the cost. This is also true for products that are exceptions to the assumptions made in the Section 7.1. This is verified as we observe for products with slow moving demand and for low cost higher type 1 service levels are obtained. The same is observed for products with fast moving demand and low cost. Likewise as the products become costly both for slow moving and fast moving products the type 1 service level goes down. For fill rate same observations are recorded.
- As expected lead time does not have a solitary impact on service levels. High type 1 service levels are obtained for products with low to medium cost. The type 1 service level drops down with cost for all three categorizations of lead time.
- The fill rate values obtained for slow moving products as expected are higher than the ones obtained from the solution. For medium and fast moving products they are very similar to the expected fill rates from the solution file.

- The variability from trial to trial in type 1 services levels and fill rate results goes down when using data generated in Minitab. There does exist though a noticeable variability especially for slow moving demand. In order to further understand this variability and to provide decision makers with the corresponding risk involved we further test our model on different finite planning horizons in Section 9.
- The error in service level measures is more impacted by change in cost and volume of demand. The error increases as the cost increases while the volume of demand decreases.

8.4 Simulation results for normally distributed demand

The simulation experiments have been carried out for demand that follows Poisson distribution. In the solution file we find safety stock under the assumptions of normal demand distribution. In this section we test our models under normal distribution to understand the error we induced when applying the model to a system governed by Poisson demand. We discretize the normally distributed demand to obtain integer values. The simulation is carried out for 27 products again. The simulation results are categorized into slow, moving and fast moving demand for simplicity. Products 1 to 9 with a mean weekly demand of 0.1 are categorized as slow moving; products 9 to 18 with mean weekly demand of 10 are medium moving, while the ones with a mean weekly demand of 100 are fast moving.

The discretization of demand is carried out in the following steps;

- The mean demand for a product i and its standard deviation are given by equation 16 in Section 5. Using these values we calculate standardized z_i score by using equation 23 below.

$\mu_i = \text{mean weekly demand for product } i.$

$\sigma_i = \text{standard deviation of mean weekly demand for product } i = \sqrt{z_i}.$

$x_i = \text{realization of normally distributed weekly demand.}$

$$z_i = \frac{(x_i - \mu_i)}{\sigma_i} \dots \dots \dots (23)$$

- After calculating the standardized z scores we calculate the cumulative distribution function Φ_i for the score using the NORMSDIST function using excel again.
- After calculating the cumulative distribution function for a mean normally distributed weekly demand we calculate the probability mass function ϕ_i by taking the following difference. The probability mass function for mean weekly demand x_i is;

$$P[D_i = x_i] = \Phi_i(x_i + 0.5) - \Phi_i(x_i - 0.5) \dots \dots \dots (24)$$

- The probability mass functions are calculated in either side till a value of zero is obtained. We finally calculate the mean and standard deviations of mean weekly demands generated to ensure they are equal to the values by having Poisson distribution.

The following table gives the discretized normal weekly demand for Poisson weekly demand with a mean of 0.1 units.

Possible demand	Probability mass function (pmf)	Cumulative distribution function (cdf)
0	0.9	0.9
1	0.09	0.99
2	0.01	1

Table 13 Discrete demand probabilities for slow moving demand.

The results obtained after carrying out simulations using normally distributed demand for medium moving products (mean weekly demand of 10) and fast moving products (mean weekly demand of 100) are similar. The results for these are presented in the Appendix C. In this section we discuss the results for slow moving products (mean weekly demand of 0.1) The corresponding results are in table 14 below, while the type 1 service levels are compared in figure 20.

Product No	Type 1 service level			Fill rate		
	Solution	Simulation (Poisson approximation)	Simulation (Normal approximation)	Solution	Simulation (Poisson approximation)	Simulation (Normal approximation)
1	0.99	1	0.99	0.99	0.97	0.96
2	0.99	0.96	0.95	0.97	0.90	0.86
3	0.85	0.75	0.76	0.39	0.69	0.65
4	0.99	0.98	0.97	0.98	0.98	0.97
5	0.98	0.98	0.96	0.97	0.94	0.90
6	0.85	0.92	0.90	0.66	0.77	0.71
7	0.99	0.99	0.97	0.99	0.99	0.97
8	0.96	0.97	0.94	0.94	0.93	0.89
9	0.85	0.82	0.77	0.80	0.80	0.74

Table 14 Comparison of Type 1 service level and Fill rates obtained from simulations under poisson and normal distribution with the expected values from solution

The results obtained by simulations under normally distributed demand are consistent with those under Poisson distribution. Very similar trends as observed in the results of the preceding section for 27 products are observed here implying change in service level measures for normally distributed demand with change in cost and nature of demand. We also obtain similar results for fill rate.




Figure 20 Comparison of Type 1 service level for slow moving products.

In Figure 21 we compare the corresponding fill rates obtained from different simulations.

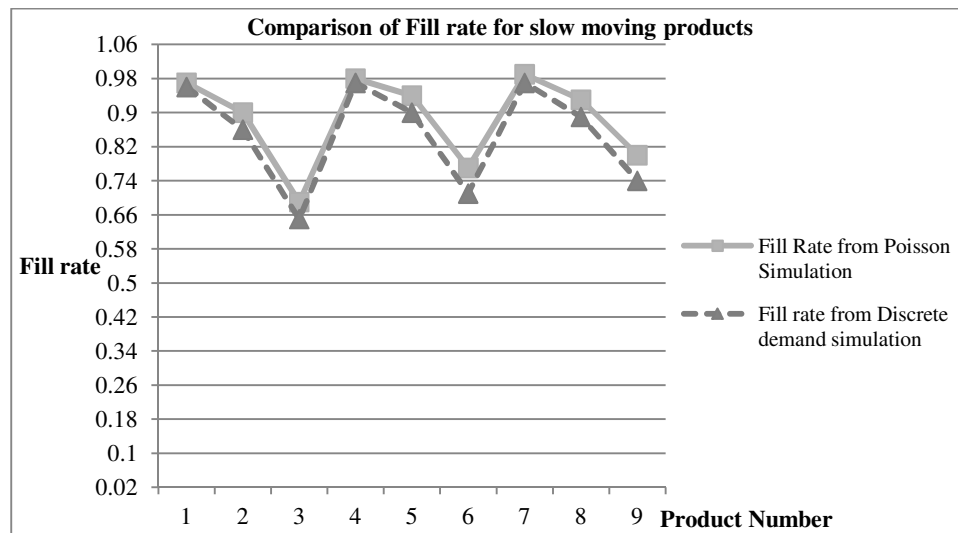


Figure 21 Comparison of Fill rates for slow moving products.

The type 1 service level values obtained after simulating demand following a Poisson distribution and demand following a normal distribution is close to each other. For slow moving products as seen in table 14 they are however not equal to the ones obtained by model solution. The difference is more noticeable in case of costly slow moving products which are product number 3, 6 and 9 in table 14 above. The results seem to follow the general trend where obtaining exact type 1 service levels and fill rates is difficult for slow moving and costlier products.

We can conclude that the model can be applied to Poisson demand confidently even though it was developed using normally distributed demand. Though to provide the decision makers about the risk and uncertainty involved for the slow moving products we carry out a rigorous risk analysis of results in the following section.

CHAPTER 9

RISK ANALYSIS

The mathematical solution to our two optimization models is verified by simulation. Simulation methods have become increasingly popular and are widely accepted in both in the academic as well as industrial community. The simulation methods though provide a realistic representation of an actual industrial set up they might provide optimal results only under specific assumptions (Banks 2007). It is necessary that the results obtained by simulation experiments have sufficient accuracy with a corresponding idea of confidence intervals (Banks 2007). If some of the input processes driving a simulation are random, runs of solution results in estimates of performance measures (Alexopoulos et al 1998). These performance estimates come with errors and it is necessary to estimate these errors. The output results hence need to be analyzed and evaluated. The reason being to understand the error involved and provide decision makers with some confidence intervals. These confidence intervals only equip the decision makers better in decision making process to achieve desired results. Notable works that emphasize validation and output analysis of simulation experiments are Carson (2002), Kleijnen (1995), Balci (2003) and Law and Kelton (2000).

In our model we assume an infinite planning horizon. The results obtained seemed to provide desired service levels under this planning horizon. To validate our models we also need to test them under finite planning horizon. There are various motivations behind this analysis. The first and most important is estimating the error in type 1 service

level values and fill rates simulation. A manufacturer may have performance contracts with its customers and also suppliers (Thomas 2005). Under these contracts failing to achieve established service levels over specified planning horizons may result in financial penalties. The manufacturers may also be a part of a larger supply chain where holding raw material and finished goods inventory may negatively impact overall system performance. The managers also end up caught in between uncertainty of demand fluctuations and strict inventory cost control carried out by higher management (Ravichandran 2007). There also might be instances where on a finite horizon planning more than required service level values are obtained. This may not be always desirable as in order to achieve these high service levels manufacturers may be incurring a high inventory cost. In validating our model we carry out simulation on finite planning horizons of one, two, four and six years. We report the corresponding variation that can be obtained in the service level values, along with estimation of quartiles. We expect that for medium and fast moving products our model to provide desired service levels on shorter planning horizons. For slow moving products we anticipate the variability in service levels to go up as the planning horizon becomes shorter and shorter.

9.1 Simulation results for planning horizon of 6 years

In this section we report the results for a planning horizon of 6 years. We increase the number of trials from 20 to 1000. There are no changes in the calculations for inventory levels and service levels, with the only change being the planning horizon consisting of 350 weeks. The first 50 weeks are considered as the warm up period weeks. In table 15 we report the average type 1 service levels and fill rates along with the minimum and maximum values that they can take. We observe a major variation in the type 1 service

level values and fill rates for slow moving products. This variation is even more noteworthy in costlier products (products 3, 6 and 9). For slow moving products the variation in fill rates is more than the variation in service levels. In our further analysis we examine products 3, 6 and 9 more closely. These are slow moving costly products and seem to have maximum variation in results.

Product No	Type 1 service level from simulation.	Type 1 Service Range, Minimum value- Maximum value.	Fill rate from simulation.	Fill rate from Range, Minimum value- Maximum value.
1	1.00	0.98-1.00	0.91	0.80-1.00
2	0.98	0.91-1.00	0.90	0.69-1.00
3	0.86	0.67-1.00	0.70	0.39-1.00
4	0.99	0.92-1.00	0.98	0.78-1.00
5	0.99	0.92-1.00	0.94	0.66-1.00
6	0.95	0.92-1.00	0.77	0.39-1.00
7	0.99	0.87-1.00	0.99	0.67-1.00
8	0.98	0.81-1.00	0.93	0.51-1.00
9	0.88	0.51-1.00	0.80	0.28-1.00
10	1.00	0.99-1.00	1.00	1.00-1.00
11	1.00	0.99-1.00	1.00	0.99-1.00
12	0.89	0.77-0.98	0.95	0.90-0.99
13	1.00	0.98-1.00	1.00	0.99-1.00
14	1.00	0.96-1.00	1.00	0.97-1.00
15	0.89	0.73-1.00	0.93	0.76-1.00
16	1.00	0.96-1.00	1.00	0.97-1.00
17	1.00	0.89-1.00	1.00	0.86-1.00
18	0.89	0.58-1.00	0.90	0.53-1.00
19	0.99	0.98-1.00	0.99	0.98-1.00
20	1.00	0.99-1.00	1.00	1.00-1.00
21	0.98	0.93-1.00	1.00	0.99-1.00
22	1.00	0.99-1.00	1.00	1.00-1.00
23	1.00	0.98-1.00	1.00	1.00-1.00
24	0.95	0.84-1.00	1.00	0.94-1.00
25	1.00	0.97-1.00	1.00	0.98-1.00
26	1.00	0.95-1.00	1.00	0.96-1.00
27	0.91	0.64-1.00	0.96	0.72-1.00

Table 15 Type 1 service levels over a planning horizon of 6 years. The average Type 1 service level and Fill rates for a 1000 trials with the minimum and maximum values.

The histograms in following figures will provide a further idea of the risk involved in using the safety stock values obtained from simulation. We first report the corresponding variation in type 1 service levels for the three products after which we report the changes in fill rates.

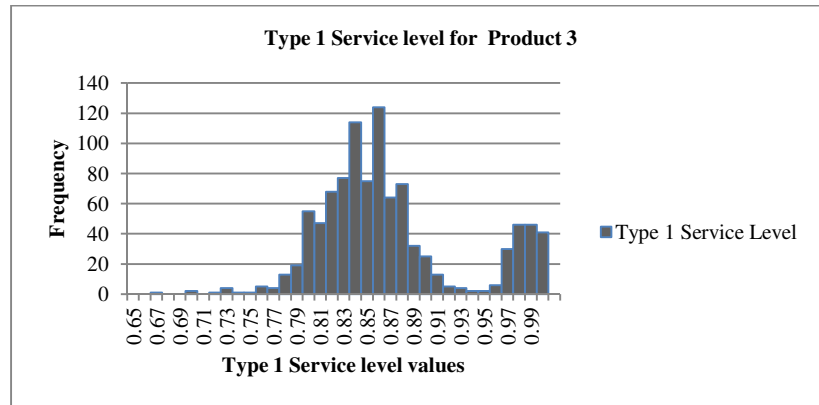


Figure 22 Variation in Type 1 service level for Product 3.

The expected type 1 service level for product 3 from the solution is 85 percent. In simulation results for a planning horizon of six years the minimum service level of near 85 percent is achieved for around seventy percent of time. The desired service level of 95 percent is achieved only 25 percent of time. In figure 23 we provide results for variation in type 1 service level for product 6, costly product with slow moving demand and medium lead time.




Figure 23 Variation in Type 1 service level for product 6.

For product 6 less variation is noticed in type 1 service level values than product 3. The minimum service level of 85 percent is achieved for about 90 percent of the iterations; while about 75 percent of the iterations give the desired service level of 95 percent is achieved.




Figure 24 Variation in Type 1 service level for Product 9.

Product 9 also exhibits significant variation in the type1 service level values. We further examine the fill rates for products 3, 6 and 9. Though for these products higher type 1 service levels are obtained for fill rates very low values are obtained. The fill rates of Product 3 and Product 6 are compared in figure 25. In both the cases around 80 percent of the time fill rate values in excess of 70 percent are observed. Consistent with the observation of type 1 service level results marginally better fill rates are observed for Product 6 than Product 3.

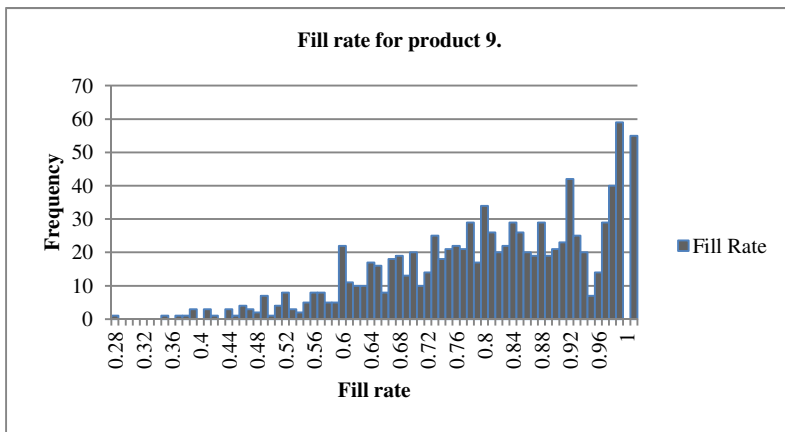
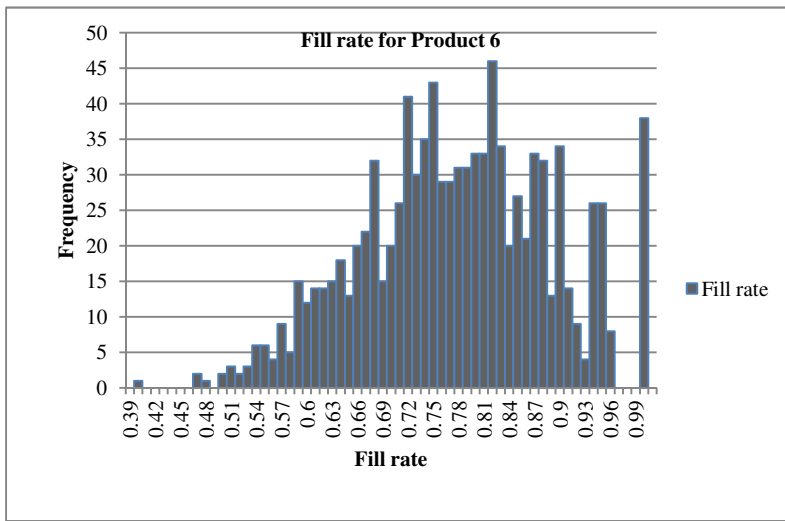
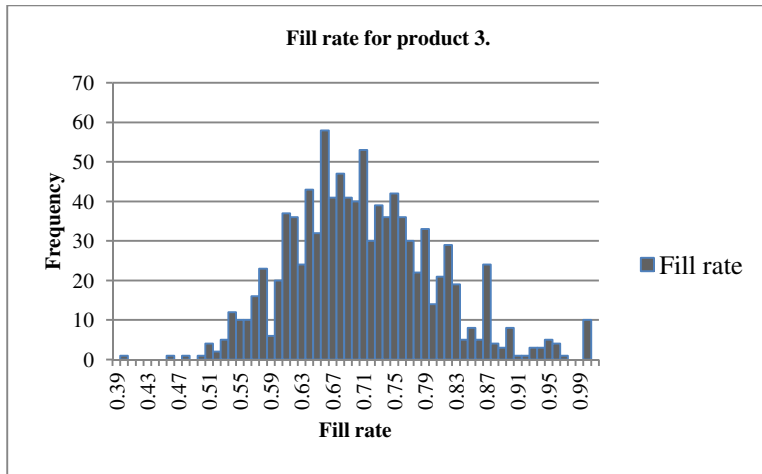


Figure 25: Variation in Fill rates for Product 3, Product 6 and Product 9.


The variation in the fill rate values is reported for Product 9. The fill rates vary significantly for Product 9 from trial to trial. On average a fill rate of around 75 percent is achieved.

To summarize the results the safety stock solutions obtained provide service levels close to the predicted model when implemented over a finite planning horizon. For slow moving products, and also for products with medium to high demand and high cost there does exist a significant variation in type 1 service levels and fill rates. The range of variation will be able to assist decision makers in assessing risk and potentially lead them to adjust their safety stock and hedge against this variation. Providing the range may not be enough especially for slow moving products. This variation is bound to increase as the planning horizon gets shorter. It becomes important to calculate the errors and develop confidence intervals for very short planning horizons. The fill rates of medium and fast moving demand are higher for a shorter planning horizon than an infinite planning horizon. This is consistent with the observations of Thomas (2005) and Chen et al (2002) who have theoretically proved that short term fill rates are higher than the long term infinite planning horizon fill rates.

In the following section we report the results for finite planning horizons of 1, 2 and 4 years. The focus is slow moving products as these seem to have maximum variation when moving from an infinite to finite planning horizon. We also provide quartile estimates and confidence intervals for all the products across the 3 planning horizons.

9.2 Simulation results for planning horizon of 1, 2 and 4 years

In this section we provide confidence interval estimates for type 1 service levels and fill rates as in the previous section. We anticipate a significant increase in the variation of results for the slow moving products. The variation is also bound to increase in case of costly medium and fast moving products as the time horizons get shorter and shorter. We first provide histograms in figures 26 and 27 for product 1, 3, 6 and 9 demonstrating the variation in type 1 service levels for shorter planning horizons.



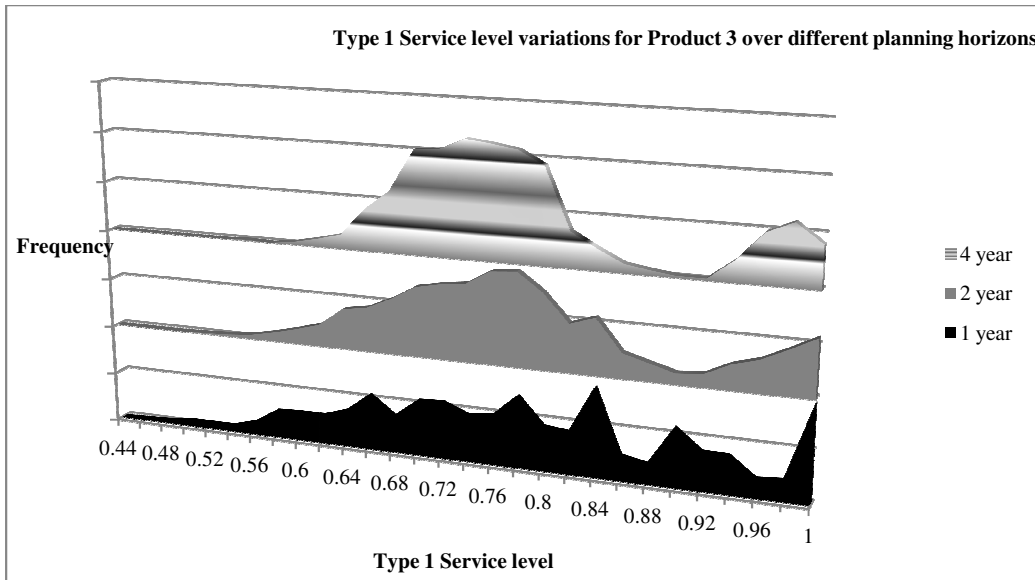



Figure 26 Type 1 service level variations for Products 1 and 3 for planning horizons of 1, 2 and 4 years.
 The variation for product 1 does not change, but it significantly increases for Product 3 especially over a planning horizon of 1 year.






Figure 27Type 1 service level variations for Products 6 and 9for planning horizons of 1, 2 and 4 years.

Similar trends are also seen for Products 3 and 9. The graphs all seem to be negatively skewed or have a long tail towards the left. The fill rate values obtained after simulation for finite planning horizons also exhibit similar variations. The corresponding histograms are attached in the Appendix C.

We now calculate confidence intervals and quartiles for all the products and across the three planning horizons. A confidence interval is the measure of error for N number of replications of a given simulation (Banks 2007). Its estimation will provide the decision makers the probability with which they can achieve desired service level measures.

While confidence interval estimation will provide the decision makers an estimate of uncertainty and error involved with the safety stock values, the quartile calculations will provide them with an idea of the risk involved. The estimation of confidence interval and

quartiles is done in Minitab using graphical data summary tool. We provide the mathematical equations here from Banks (2007) which are used in the background.

We define the following variables for estimating confidence intervals.

$Y_i =$ Type 1 Service level or Fill rate for replication i .

$R =$ number of replications for simulation,

$\bar{Y} =$ mean type 1 Service level for R replications, calculated by equation (25).

$$\bar{Y} = \sum_{i=1}^R Y_i \dots \dots \dots (25)$$

$S^2 =$ variance across the R replications, given by equation (26).

$$S^2 = \frac{1}{R-1} \sum_{i=1}^R (Y_i - \bar{Y})^2 \dots (26)$$

$t_{(\frac{\alpha}{2}, R-1)} =$ quantile for confidence interval of $(1 - \alpha)$ and $R - 1$ replications.

The confidence interval for a mean of \bar{Y} is given by the following equation (27),

$$\bar{Y} \pm t_{(\frac{\alpha}{2}, R-1)} \frac{S}{\sqrt{R}} \dots \dots \dots (27)$$

The results for confidence intervals for a planning horizon of 4 years are provided in table 16 below. We observe that variation is higher in the slow moving products then the medium moving and fast moving products. Hence the results for the first 9 products are provided here and the other results are provided in Appendix C.

Product No	Type 1 service level from simulation. Mean Service Level \bar{Y}	Confidence interval of 95 percent.	Margin of error for 95 percent confident interval.	Minimum Type 1 service level	Maximum Type 1 service level
1	0.9912	(0.9910-0.9914)	0.0116	0.962	0.999
2	0.9630	(0.9628-0.9640)	0.1934	0.881	0.998
3	0.7780	(0.7710-0.7840)	0.1937	0.557	0.999

4	1.0000	(0.9991-1.0000)	0.0182	0.938	1.000
5	0.9850	(0.9849-0.9870)	0.0451	0.876	1.000
6	0.9225	(0.9190-0.9261)	0.1013	0.727	1.000
7	0.9690	(0.9659-0.9720)	0.1030	0.694	0.999
8	0.9643	(0.9612-0.9673)	0.1025	0.690	1.000
9	0.8112	(0.8032-0.8190)	0.2499	0.333	0.999

Table 16: 95 percent confidence intervals, minimum and maximum Type 1 service level values for slow moving products for a planning horizon of 4 years.

We calculate a 95 percent confidence interval for the mean type 1 service level values obtained for a planning horizon of 4 years and 1000 replications. The second column in the above table gives the 95 percent confidence interval limits for the mean given in the first column. The results can be interpreted as for Product 9 the probability that type 1 service level is between 0.803 and 0.819 is 95. We also provide the minimum values obtained for type 1 service level after 1000 replications. Though the minimum values are small the confidence intervals are small indicating a high probability of achieving desired type 1 service levels. This implies that the model provides good results for a planning horizon of 4 years.

In figure 28 a graphical summary of quartile calculation is provided for the slow moving products. The number of quartiles for a given set of data is usually 4. We choose this representation as providing the minimum value is important in order to understand the risk associated with using a particular safety stock. The 5 quartiles calculated are as follows;

- 0 or the 1st Quartile represents the minimum value obtained for 1000 replications.
- 2nd Quartile represents the type 1 service level value which is higher or equal to the lower 25 percent of the type 1 service level values obtained after 1000 replications.

- 3rd Quartile represents the type 1 service level value which is higher or equal to the lower 50 percent of the type 1 service level values obtained after 1000 replications.
- 4th Quartile represents the type 1 service level value which is higher or equal to the lower 75 percent of the type 1 service level values obtained after 1000 replications.
- 5th Quartile provides the maximum value obtained for type 1 service level after a 1000 replications.




Figure 28 Change in Type 1 service level Quartiles for slow moving products for a planning horizon of 4 years.

The minimum values and the first quartiles for type 1 service level are very low as compared to the minimum type 1 service level of 85 percent. However there is a sharp increase in the quartile values. This indicates that about 25 percent of the replications provide low type 1 service levels. Based on the quartiles provided and the minimum and maximum type 1 service level values obtained the decision makers can adjust their safety stock values to achieve desired service levels.

Product No	Fill rate from simulation. Mean Service Level(\bar{Y})	Confidence interval of 95 percent.	Margin of error for 95 percent confident interval	Minimum Fill rate.	Maximum Fill rate.
1	0.9772	(0.9752-0.9793)	0.0679	0.800	0.999
2	0.9130	(0.9100-0.9164)	0.1855	0.711	0.999
3	0.7042	(0.6981-0.7100)	0.1836	0.412	0.999
4	0.9860	(0.9840-0.9882)	0.0620	0.803	1.000
5	0.9469	(0.9421-0.9479)	0.1168	0.924	0.999
6	0.7721	(0.7654-0.7762)	0.2223	0.395	1.000
7	0.9462	(0.9412-0.9513)	0.1519	0.556	0.999
8	0.9560	(0.9411-0.9611)	0.1519	0.557	1.000
9	0.8042	(0.7951-0.8134)	0.2856	0.309	0.999

Table 17: 95 percent confidence intervals; minimum and maximum Fill rate values for slow moving products for a planning horizon of 4 years.

Fill rates are a better indicator of service level measures on shorter planning horizons for slow moving products. It is crucial that desired fill rates are achieved on a weekly basis along with obtaining desired type 1 service levels for planning horizons. We computed both the average weekly fill rate and fill rates across the entire planning horizons. While calculating these values we segregate the weeks where no demand is realized and only account for products where demand is to be satisfied. We observed that average weekly fill rates are again equal to the fill rates for entire planning horizons. We report fill rates for the entire planning horizon in our analysis. Similar results and calculations are also carried out for fill rates for slow moving products and across a planning horizon of 4 years. The corresponding results are in table 17 above and figure 29 below. The fill rate quartiles exhibit similar characteristics as those of type 1 service level values. For Products 1, 2,4,5,7 and 8 the probability of achieving fill rates close to the mean values is high with small confidence interval. As the type 1 service level values and fill rates are

both high we conclude that the safety stock solution results work for a shorter planning horizon of 4 years.




Figure 29 Change in Fill rate Quartiles for slow moving products for a planning horizon of 4 years.

For products 3, 6 and 9 though fill rates below the desired levels are obtained, the confidence intervals are small again. The small confidence intervals indicate small error in the simulation results. The quartile calculations of fill rates are very similar to that of type 1 service levels. Products 3, 6 and 9 which are slow moving and costly products 1st and 2nd Quartiles are on the lower side. For other products only the 1st quartiles have low value, indicating at in 75 percent of the replications minimum desired fill rates are achieved.

We now estimate the confidence intervals and quartiles for planning horizons of 1 and 2 years. The aim is to test the probability with which desired type 1 service levels and fill rates will be achieved for shorter planning horizons. In the following table 18 we look at

confidence intervals of type 1 service levels for all the products across a planning horizon of 2 years.

Product No	Type 1 service level from simulation. Mean Service Level (\bar{Y})	Confidence interval of 95 percent.	Margin of error for 95 percent confident interval	Minimum Type 1 service level.	Maximum Type 1 service level.
1	0.9962	(0.9951-0.9965)	0.0173	0.912	0.999
2	0.9640	(0.9630-0.9660)	0.0549	0.899	1.000
3	0.7782	(0.7712-0.7850)	0.2146	0.524	1.000
4	0.9860	(0.9852-0.9870)	0.0266	0.846	0.990
5	0.9701	(0.9682-0.9733)	0.0669	0.682	0.990
6	0.9092	(0.9052-0.9142)	0.1360	0.586	0.971
7	0.9861	(0.9840-0.9870)	0.0481	0.740	0.990
8	0.9462	(0.9411-0.9502)	0.1404	0.504	0.999
9	0.7951	(0.7851-0.8062)	0.3332	0.133	0.980
10	0.9804	(0.9802-0.9810)	0.0026	0.961	0.980
11	0.9801	(0.9793-0.9803)	0.0061	0.952	0.980
12	0.8621	(0.8592-0.8641)	0.0893	0.838	0.987
13	0.9802	(0.9803-0.9812)	0.0050	0.933	0.980
14	0.9794	(0.9784-0.9799)	0.0158	0.876	0.990
15	0.8352	(0.8302-0.8419)	0.1699	0.495	0.980
16	0.9804	(0.9801-0.9815)	0.0056	0.914	0.980
17	0.9781	(0.9782-0.9791)	0.0252	0.828	0.980
18	0.8442	(0.8353-0.8533)	0.2953	0.106	1.000
19	0.9809	(0.9809-0.9819)	0.0008	0.971	0.980
20	0.9804	(0.9803-0.9815)	0.0029	0.961	0.988
21	0.9572	(0.9562-0.9582)	0.0374	0.876	0.989
22	0.9804	(0.9809-0.9908)	0.0005	0.971	0.999
23	0.9808	(0.9806-0.9809)	0.0051	0.942	0.980
24	0.9311	(0.9284-0.9348)	0.0969	0.685	0.990
25	0.9868	(0.9837-0.9869)	0.0000	0.912	0.980
26	0.9804	(0.9804-0.9818)	0.0070	0.914	0.980
27	0.8869	(0.8799-0.8938)	0.2263	0.361	0.980

Table 18 95 percent confidence intervals; minimum and maximum Type 1 service level values for all products for a planning horizon of 2 years.

The table shows that the 95 percent confidence intervals increase for all products. Products with medium moving demand and high cost, which are Products no 12, 15 and 18, also exhibit a large variation in the results. Similar characteristics are observed for fast moving costly products. For products with slow, medium and fast moving demand but with low to medium cost the confidence intervals continue to be small ensuring the applicability of the model solutions to a finite planning horizon of 2 years. We highlight the products with maximum variation between their minimum and maximum type 1 service level values and provide change in quartile values for them in the following figure 30.




Figure 30 Change in Type 1 service level Quartiles for highlighted products in Table 18 for a planning horizon of 2 years.

The results for a planning horizon of 2 years are similar to the planning horizon of 4 years. The minimum type 1 service level values and the 1st quartiles are lower than the planning horizon of 4 years. The lines though are steeper indicating a sharp increase in the 2nd Quartile values. We can imply from Figure 30 that near optimal type 1 service levels are obtained for a planning horizon of 2 years for about 75 percent of the time. In

the following table and figure we look at confidence intervals and quartiles of fill rates for planning horizon of 27 products.

Product No	Fill rate from simulation. Mean Service Level (\bar{Y})	Confidence interval of 95 percent.	Margin of error for 95 percent confident interval	Minimum Fill rate.	Maximum Fill rate.
1	0.9781	(0.9752-0.9819)	0.0899	0.553	1.000
2	0.9192	(0.9151-0.9238)	0.1295	0.420	1.000
3	0.7199	(0.7114-0.7264)	0.2395	0.182	1.000
4	0.9888	(0.9852-0.9903)	0.0837	0.570	1.000
5	0.9526	(0.9468-0.9572)	0.1540	0.470	1.000
6	0.7895	(0.7794-0.7981)	0.3058	0.020	1.000
7	0.9924	(0.9891-0.9941)	0.0803	0.425	1.000
8	0.9467	(0.9411-0.9514)	0.1519	0.548	1.000
9	0.8266	(0.8149-0.8384)	0.3645	0.050	1.000
10	0.9997	(0.9990-0.9999)	0.0008	0.990	1.000
11	0.9993	(0.9990-1.0000)	0.0020	0.983	1.000
12	0.9694	(0.9684-0.9702)	0.0340	0.886	1.000
13	0.9994	(0.9990-1.0000)	0.0013	0.980	1.000
14	0.9998	(0.9997-1.000)	0.0068	0.947	1.000
15	0.9257	(0.9211-0.9299)	0.1127	0.689	1.000
16	0.9999	(0.9990-1.0000)	0.0015	0.979	1.000
17	0.9999	(0.9981-0.9999)	0.0120	0.878	1.000
18	0.9021	(0.8951-0.9091)	0.2323	0.217	1.000
19	1.0000	(0.9990-1.0000)	6.08E-05	0.993	1.000
20	1.0000	(0.9997-1.0000)	0.0002	0.998	1.000
21	0.9987	(0.9942-0.9989)	0.00352	0.998	1.000
22	1.0000	(0.9999-1.0000)	4.51E-05	0.989	1.000
23	0.9999	(0.9990-0.9999)	0.0004	0.999	1.000
24	0.9934	(0.9931-0.9941)	0.0175	0.996	1.000
25	0.9994	(0.9982-0.9999)	0.000	0.938	1.000
26	0.9992	(0.9991-1.0000)	0.0011	0.984	1.000
27	0.9751	(0.9721-0.9772)	0.0865	0.699	1.000

Table 19 95 percent confidence intervals; minimum and maximum Type 1 service level values for all products for a planning horizon of 2 years.

The same products that exhibit a large variation in type 1 service level also exhibit a large variation in fill rate values for the 4 year planning horizon. The minimum fill rate values obtained are also extremely low for these products. We again highlight products with large variation and estimate the corresponding quartiles. The products with large variation in fill rates are all the slow moving products and medium and fast moving products with high cost. We divide the quartile estimation in two parts in figure 31 we look at quartiles for only slow moving products and figure 32 we look at quartiles for medium and fast moving products with high cost.




Figure 31 Change in Fill rate Quartiles for slow moving products (1-9) for a planning horizon of 2 years.

The quartile estimates for a planning horizon of 2 years are very similar to a planning horizon of 4 years. For slow moving products above we have very low minimum and 1st quartile values though they significantly increase as indicated by steep graph from the 2nd quartile values. For products 1, 2, 4,5,7,8 desired fill rates seem to be obtained in the 2nd quartile itself. Similar trends are observed for Products 12, 15, 18 and 27. This ensures applicability of the model solutions for a planning horizon of 2 years.

The results for type 1 service levels for a planning horizon of 1 year are provided in table 20 below.

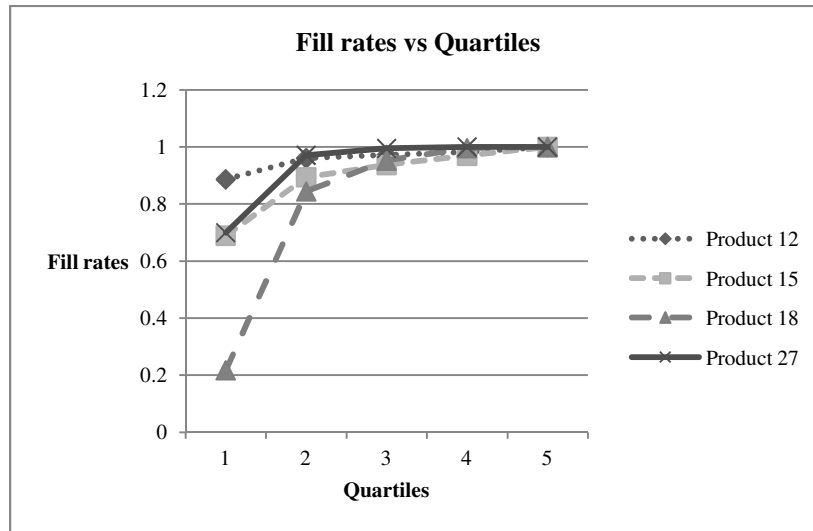


Figure 32 Change in Fill rate Quartiles for highlighted products in Table 19 for a planning horizon of 2 years.

Product No	Type 1 service level from simulation. Mean Service Level (\bar{Y})	Confidence interval of 95 percent.	Margin of error for 95 percent confident interval	Minimum Type 1 service level.	Maximum Type 1 service level.
1	0.9962	(0.9952-0.9968)	0.0208	0.940	1.000
2	0.9640	(0.9614-0.9669)	0.0771	0.790	1.000
3	0.7850	(0.7778-0.7934)	0.2495	0.452	1.000
4	0.9961	(0.9952-0.9974)	0.0336	0.830	1.000
5	0.9809	(0.9778-0.9832)	0.0889	0.735	1.000
6	0.9224	(0.9164-0.9289)	0.1864	0.563	1.000
7	0.9961	(0.9944-0.9972)	0.0482	0.662	1.000
8	0.9652	(0.9591-0.9718)	0.1851	0.281	1.000
9	0.8208	(0.8070-0.8360)	0.4479	0.010	1.000
10	0.9999	(0.9999-1.0000)	0.0020	0.981	1.000
11	0.9999	(0.9989-0.9999)	0.0078	0.960	1.000
12	0.8800	(0.8762-0.8840)	0.1260	0.644	1.000
13	0.9999	(0.9999-1.0000)	0.0088	0.867	1.000

14	0.9978	(0.9974-0.9989)	0.0228	0.830	1.000
15	0.8505	(0.8420-0.8572)	0.2392	0.283	1.000
16	0.9999	(0.9990-0.9999)	0.0145	0.811	1.000
17	0.9969	(0.9944-0.9987)	0.0525	0.622	1.000
18	0.8504	(0.8384-0.8623)	0.3777	0.754	1.000
19	0.9999	(0.9999-1.0000)	0.0026	0.962	1.000
20	0.9999	(0.9999-1.0000)	0.0057	0.962	1.000
21	0.9739	(0.9721-0.9754)	0.0593	0.830	1.000
22	0.9999	(0.9999-1.0000)	0.0057	0.962	1.000
23	0.9999	(0.9999-1.0000)	0.0026	0.962	1.000
24	0.9404	(0.9351-0.9459)	0.1532	0.528	1.000
25	0.9999	(0.9999-1.0000)	0.0023	0.962	1.000
26	0.9999	(0.9999-1.0000)	0.0141	0.792	1.000
27	0.8834	(0.8872-0.8949)	0.3320	0.050	1.000

Table 20 95 percent confidence intervals; minimum and maximum Type 1 service level values for all products for a planning horizon of 1 years.

The variation in type 1 service level widens across the planning horizon of 1 year. The products with medium demand and high demand continue to have low confidence intervals indicating that the solution to the model provides more desired service level measures for a planning horizon of 1 year as well. The slow moving products exhibit a wide variation and estimate the quartiles along with the highlighted products again to estimate the risk involved in it.




Figure 33Change in Type 1 service level Quartiles for slow moving products for a planning horizon of 1 year.

The quartiles for slow moving products for a planning horizon of 1 year are represented above. As similar results are achieved for a planning horizon we conclude the model solutions also work for finite planning horizons. The quartile estimation for other highlighted products is similar to slow moving products where desired type 1 service levels are achieved after the 1st quartile. They are represented in figure 34 below.




Figure 34Change in Type 1 service level Quartiles for highlighted products in Table 20for a planning horizon of 1 year.

The fill rate estimates on a planning horizon of 1 year are very crucial especially in the case of slowing moving products. There will be demand of one to two odd products and it is important that 100 percent fill rate is achieved whenever demand is realized. We look at fill rate estimates for a planning horizon of 1 year in the following table 21.

Product No	Fill rate from simulation. Mean Service Level (\bar{Y})	Confidence interval of 95 percent.	Margin of error for 95 percent confident interval	Minimum Fill rate.	Maximum Fill rate.
1	0.9829	(0.9791-0.9862)	0.1025	0.685	1.000
2	0.9274	(0.9221-0.9334)	0.1629	0.554	1.000
3	0.7342	(0.7240-0.7749)	0.3195	0.247	1.000
4	0.9892	(0.9872-0.9929)	0.0843	0.585	1.000
5	0.9591	(0.9539-0.9649)	0.1800	0.462	1.000
6	0.8229	(0.8102-0.8342)	0.3817	0.030	1.000
7	0.9924	(0.9902-0.9959)	0.0853	0.500	1.000
8	0.9600	(0.9538-0.9667)	0.2157	0.150	1.000
9	0.8554	(0.8414-0.8689)	0.4216	0.050	1.000
10	0.9995	(0.9990-0.9999)	0.0003	0.999	1.000
11	0.9999	(0.9998-1.0000)	0.0016	0.860	1.000
12	0.9694	(0.9680-0.9710)	0.0450	0.875	1.000
13	0.9999	(0.9999-1.000)	0.0044	0.929	1.000
14	0.9999	(0.9998-0.999)	0.0106	0.853	1.000
15	0.9244	(0.9199-0.9290)	0.1548	0.499	1.000
16	0.9905	(0.9989-0.9999)	0.0041	0.950	1.000
17	0.9988	(0.9967-0.9999)	0.0319	0.691	1.000
18	0.9056	(0.8961-0.9150)	0.2950	0.120	1.000
19	0.9999	(0.9999-1.0000)	9.02E-05	0.998	1.000
20	0.9999	(0.9989-1.0000)	0.0003	0.994	1.000
21	0.9987	(0.9967-0.9989)	0.0059	0.969	1.000
22	0.9999	(0.9989-0.9999)	0.0003	0.994	1.000
23	1.0000	(0.9949-1.0000)	0.0002	0.996	1.000
24	0.9927	(0.9921-0.9929)	0.0294	0.886	1.000
25	1.0000	(0.9999-1.0000)	0.0002	0.995	1.000
26	0.9999	(0.9989-1.0000)	0.0035	0.945	1.000
27	0.9704	(0.9662-0.9739)	0.1226	0.420	1.000

Table 21 95 percent confidence intervals; minimum and maximum Fill rate values for all products for a planning horizon of 1 years.

The slow moving and costly products again exhibit significant variation in fill rate values.

The other products have low confidence intervals and less variation.

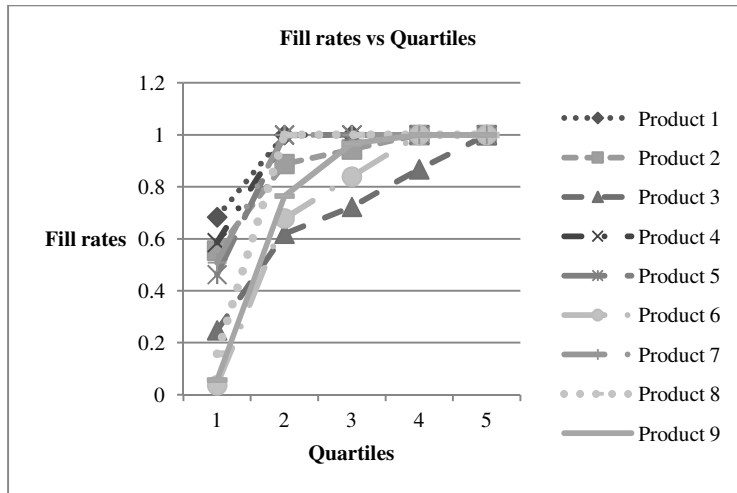


Figure 35 Change in Fill rate Quartiles for slow moving products for a planning horizon of 1 year.

We see a similar graph as that of fill rate quartiles for a planning horizon of 2 and 4 years for the slow moving products. The fill rate quartiles after the 1st quartile rise significantly. The type 1 service level and fill rates have narrow confidence intervals and show improvement in the quartile values after the first quartile for slow moving products. We can conclude from the above figure that other highlighted products will also have similar quartile measures. We can safely assume that for a planning horizon of 1 year as well the two models provide good implementable solutions.

The models provide solutions with performance close to that predicted for shorter planning horizons as well. The narrow confidence intervals across all planning horizons imply a high probability of achieving desired service level measures. These confidence intervals consistently remain narrow across all the planning horizons for products with medium to high demand and cost. The model can hence safely be applied to inventory of such products. The slow demand and costly products exhibit variation in results across all the different shorter planning horizons. The quartiles of these slow moving products

though improve after the 1st quartile; thus there exists around 25 percent of a chance of having low service levels. These products with a high variation and low first quartiles can be categorized as sensitive products. The decision makers can adjust the service level measures for these sensitive products. The tradeoff between inventory cost and desired service levels along with estimates of confidence intervals should help decision makers balance their inventory level and inventory cost.

CHAPTER 10

LIMITATIONS

The fill rate equation in the two models can lead to double counting backorders which as seen in the previous sections leads to fill rate predictions that heavily underestimate the actual realizations. In our simulation we have eliminated this problem by accounting for inventory first by deducting the backorders and then estimating the service level measures. Similar efforts can be taken while implementing the model in an actual industrial setup.

Our work is motivated by spare parts inventory and we have assumed that demand for these parts or products are not related. In an actual industrial setup the demand of one part generates demand for other minor parts. A simple example can be demand for an oil pump which will also generate demand for oil gaskets and corresponding nuts, bolts and screws used in the oil pump assembly process. This connection can at times make it difficult to achieve desired service level measures. Our work can provide a starting reference point in such cases. The products can be grouped together in clusters. The products then can be classified as primary products and assembly products. Decision makers can first decide on the safety stock for the primary products then estimate the demand of subsequent assembly products and consequently decide the safety stock of the assembly products. Our model can also be used as labeling these clusters as products and then optimizing safety stocks for these clusters. The individual demands can be then

separated understanding the connection between the different parts and individual safety stocks can then be decided.

The models developed are for a single echelon supply chain. A supply chain also has desired service level targets. In such cases safety stock and inventory levels are connected across different echelons. Supply chains also include members from third party logistics where correct information sharing may not always happen. In these situations obtaining a high service level across one echelon may negatively impact the other partners in the supply chain. The results from our models can provide an initiation step for heuristic solutions which can be developed for optimizing service level measures for an entire supply chain. The different partners can also decide different demand points in the supply chain and optimize their safety stock values towards meeting this demand in time.

We have also assumed a constant lead time while carrying out the simulation. In an actual industrial setup the delivery of an item may not always occur on time. In such cases the safety stock values can be adjusted to hedge against the uncertainty in lead time. We however believe that optimizing under such situations will only lead to overestimation of safety stock. This will also increase the inventory cost. By using our model decision makers can adjust their priorities and accordingly adjust the safety stock values.

CHAPTER 11

CONCLUSION

In our work we have provided an inventory control tool for optimizing the safety stock values of spare parts. The model is computationally simple as by changing the Lagrangian multiplier the changes in desired service level measures can be observed. The model provides optimal service level measures for a wide variety of products and across different planning horizons. The results by mathematical solution provide desired individual service levels as well as aggregate service level measures. Our results were validated using discrete event simulation in Excel. The simulations were carried out for Poisson demand. We also verified our results for normalized discrete demand. The simulation across different planning horizons was carried out to estimate the error between safety stock values obtained from solution to the values that practitioners using this tool can experience in the short term, when they evaluate their inventory performance. The extremely low confidence intervals and high 2nd quartile values imply a low error and very little risk in applying these models.

The closed form solution suggests the safety stock levels are impacted by change in volume of demand, lead time and cost. We observed that with extremely slow products achieving desired service level measures is difficult for costly products. For slow moving products with low to medium cost the desired service level measures are easily obtained. For products with medium to high volume desired service levels are obtained on infinite planning horizon. As the planning horizons get shorter the variability and risk for all

costly products increases. We provide confidence intervals and quartiles in these short planning horizons which will help understand this uncertainty. The inventory simulated also contained different products and consistent results were obtained for all the products.

The work can be extended using a better equation for fill rate calculation. There can be additional constraints added for modeling the connection between different products for inventories of assembled products. Though we generated different types of products for simulation purpose the inventories followed general assumptions of industrial settings. Under these assumptions products with slow or low volume of demand were considered costly while the products with fast or high volume of demand were considered as cheaper. We included products that were exception to the above rule however the total number of such products in the inventory was low. The model can be tested for different types of inventories where such products are high or equal in number to products that follow generalized market guidelines. It would also be useful to test the robustness of our model's safety stock solutions for demands other than Poisson and Normal. Finally our models can be extended to multi echelon systems.

APPENDIX A

LAGRANGIAN SOLUTIONS OF THE FIRST MODEL

The first model is

$$\text{Minimize cost: } \sum_{i=1}^N c_i z_i \sigma_{LTD} \dots \dots \dots (16)$$

Subject to

$$\Phi(z_i) \geq MSL^i \text{ for } i = 1 \dots N \dots \dots \dots (17)$$

$$\sum_{i=1}^N \frac{\mu_i}{\mu_{total}} \Phi(z_i) \geq ASL \dots \dots \dots (18)$$

Let λ = lagrangian multiplier introduced for the second constraint.

We apply lagrangian relaxation to the equation 18. Before we carry out the lagrangian differentiation the corresponding term is multiplied by $\frac{1}{C}$, where C is the total inventory cost of all the components. This is done to make sure that the solution obtained for λ , after carrying out the lagrangian multiplier is between 0 and 1. This will eventually provide us with closed form solutions for λ . This is done for both the models.

The lagrangian relaxation derivation is carried out as follows;

$$\frac{1}{C} \sum_{i=1}^N c_i z_i \sigma_{ltd} - \lambda \left(\left(\sum_{i=1}^N \frac{\mu_i}{\mu_{total}} \Phi(z_i) \right) - ASL \right) \dots \dots \dots (19)$$

In equation 19 $\Phi(z_i)$ is the cumulative distribution function of normal demand for product i.

$$\frac{1}{C} \sum_{i=1}^N c_i z_i \sigma_{ltd} - \lambda \left(\left(\sum_{i=1}^N \frac{\mu_i}{\mu_{total}} \Phi(z_i) \right) - ASL \right) \dots \dots \dots (20)$$

We can carry out the differentiation as;

$$\frac{d}{dz_i} \left\{ \frac{1}{C} \sum_{i=1}^N c_i z_i \sigma_{ltd} - \lambda \left(\left(\sum_{i=1}^N \frac{\mu_i}{\mu_{total}} \Phi(z_i) \right) - ASL \right) \right\} =$$

The differentiation is carried out for every product i, which eliminate the summation sign above;

$$\frac{d}{dz_i} \left\{ \frac{c_i z_i \sigma_{LTD}}{C} - \lambda \left(\frac{\mu_i}{\mu_{Total}} \Phi(z_i) \right) \right\} =$$

Carrying out the differentiation, the derivate of cumulative distribution function for normal distribution is the probability distribution function represented by (z_i);

$$\left\{ \frac{c_i \sigma_{LTD}}{C} - \lambda \left(\frac{\mu_i}{\mu_{Total}} \phi(z_i) \right) \right\} =$$

For normal distribution (z_i) = $\frac{1}{\sqrt{2\pi}} e^{-\frac{z_i^2}{2}}$

$$\left\{ \frac{c_i \sigma_{LTD}}{C} - \lambda \left(\frac{\mu_i}{\mu_{Total}} \frac{1}{\sqrt{2\pi}} e^{-\frac{z_i^2}{2}} \right) \right\} =$$

Solving the above equation for z_i;

$$\left\{ \frac{c_i \sigma_{LTD}}{C} = \lambda \left(\frac{\mu_i}{\mu_{Total}} \frac{1}{\sqrt{2\pi}} e^{-\frac{z_i^2}{2}} \right) \right\}$$

$$\frac{\sqrt{2\pi} c_i \sigma_{LTD} \mu_{Total}}{\lambda C \mu_i} = e^{-\frac{z_i^2}{2}} \dots \dots \dots (21)$$

Taking natural logarithms on both sides of the above equation;

$$\ln \frac{\sqrt{2\pi} c_i \sigma_{LTD} \mu_{Total}}{\lambda C \mu_i} = \frac{-z_i^2}{2}$$

Carrying out further simplification the closed form solution for z_i as follows;

$$z_i = \sqrt{-2 \ln \left[\frac{([\sqrt{2}])c_i \sigma_{LTD} \mu_i}{\lambda C \mu_i} \right]} \dots \dots \dots (22)$$

There are two z_i values obtained. One obtained by constraint 17 and the other that we get from equation 22 above. The maximum value of the two is chosen to ensure that the minimum Type 1 service level is achieved along with achieving an aggregate system service level.

The closed form solution of the model above is;

$$z_i = \text{Max} \left[\sqrt{-2 \ln \left[\frac{(\sqrt{2\pi})[c_i \sigma_{LTD} \mu_{Total}]}{C \lambda \mu_i} \right]}, \Phi^{-1}(MSL^i) \right] \dots \dots \dots (23)$$

APPENDIX B

LAGRANGIAN SOLUTIONS OF THE SECOND MODEL

The second model is;

$$\text{Minimize cost: } \sum_{i=1}^N c_i z_i \sigma_{LTD} \dots (24)$$

Subject to

$$(fr_i(z_i)) \geq MFR^i \quad \text{for } i = 1 \dots N \dots (25)$$

$$\sum_{i=1}^N \frac{\mu_i}{\mu_{total}} fr_i(z_i) \geq FR \dots (26)$$

For the second constraint we write Fill rate $fr_i(z_i)$ using $L(z_i)$. The general expression for $L(z_i)$ is,

$$\begin{aligned} L(z) &= \left(\int_z^\infty (t - z) \phi(t) dt \right) = \\ &= \int_z^\infty t \phi(t) dt - z(1 - \Phi(z)) = \\ &= \Phi(z) - z(1 - \Phi(z)) \dots (27). \end{aligned}$$

For every product i the expression above will change to,

$$L(z_i) = \Phi(z_i) - z_i(1 - \Phi(z_i)) \dots (28)$$

The expression that connects $fr_i(z_i)$ and $L(z_i)$ is,

$$L(z_i) = (1 - fr_i(z_i)) \frac{\mu_i}{\sigma_{LTD}^i} \dots (29)$$

Combining equations 28 and 29 we get,

$$f r_i(z_i) = \left[1 - \frac{\{(\phi(z_i) - z_i (1 - \Phi(z_i))) \sigma_{LTD}^i\}}{\mu_i} \right] \dots (30)$$

Substituting it in the second constraint we rewrite the equation as;

$$\sum_{i=1}^N \frac{\mu_i}{\mu_{total}} \left[1 - \frac{\{(\phi(z_i) - z_i (1 - \Phi(z_i))) \sigma_{LTD}^i\}}{\mu_i} \right] \geq FR \dots \dots (31)$$

Taking the Lagrangian relaxation for the second constraint;

$$\frac{1}{C} \sum_{i=1}^N c_i z_i \sigma_{LTD}^i - \lambda \left[\left(\sum_{i=1}^N \frac{\mu_i}{\mu_{Total}} \left[1 - \frac{\{(\phi(z_i) - z_i (1 - \Phi(z_i))) \sigma_{LTD}^i\}}{\mu_i} \right] \right) - FR \right]$$

Taking the derivatives;

$$\frac{d}{dz_i} \left\{ \frac{1}{C} c_i z_i \sigma_{LTD}^i - \lambda \left[\left(\frac{\mu_i}{\mu_{Total}} \left[1 - \frac{\{(\phi(z_i) - z_i (1 - \Phi(z_i))) \sigma_{LTD}^i\}}{\mu_i} \right] \right) - FR \right] \right\}$$

The differentiation for different parts is;

$$\frac{d}{dz_i} \left\{ \frac{1}{C} c_i z_i \sigma_{LTD}^i \right\} = \frac{1}{C} c_i \sigma_{LTD}^i \dots \dots \dots (32)$$

For the second part of the equation we have;

$$\frac{d}{dz_i} \left\{ -\lambda \left[\left(\frac{\mu_i}{\mu_{Total}} \left[1 - \frac{\{(\phi(z_i) - z_i (1 - \Phi(z_i))) \sigma_{LTD}^i\}}{\mu_i} \right] \right) \right] \right\}$$

Separating the constants and taking the derivatives,

$$-\frac{\lambda}{\mu_{Total}} \frac{d}{dz_i} (\mu_i - (\phi(z_i) - z_i (1 - \Phi(z_i))) \sigma_{LTD}^i)$$

Sub for $\phi(z_i)$ the probability distribution function and taking derivate of one;

$$-\frac{\sigma_{LTD}^i \lambda}{\mu_{Total}} \frac{d}{dz_i} \left\{ - \left(\frac{1}{\sqrt{2\pi}} e^{-\frac{z_i^2}{2}} - z_i (1 - \Phi(z_i)) \right) \right\} =$$

Carrying out the differentiation,

$$-\frac{\sigma_{LTD}^i \lambda}{\mu_{Total}} [-\{-z_i \Phi(z_i) - 1 + z_i \Phi(z_i) + \Phi(z_i)\}]$$

This further simplifies too;

$$-\frac{\sigma_{LTD}^i \lambda}{\mu_{Total}} \{-[-1 + \Phi(z_i)]\} \dots \dots \dots (33)$$

or;

$$-\frac{\sigma_{LTD}^i \lambda}{\mu_{Total}} \{1 - \Phi(z_i)\} \dots \dots \dots (34)$$

The last term FR is constant and hence the derivate of the last term is zero. Combining the results obtained by equation 32 and 34;

$$\frac{1}{C} c_i \sigma_{LTD}^i - \frac{\sigma_{LTD}^i \lambda}{\mu_{Total}} \{1 - \Phi(z_i)\} = 0 \dots \dots \dots (35)$$

Solving for $\Phi(z_i)$ the cumulative distribution function for z_i ,

$$-\frac{\sigma_{LTD}^i \lambda}{\mu_{Total}} \{1 - \Phi(z_i)\} = -\frac{1}{C} c_i \sigma_{LTD}^i$$

$$\frac{\sigma_{LTD}^i \lambda}{\mu_{Total}} \{\Phi(z_i)\} = -\frac{1}{C} c_i \sigma_{LTD}^i + \frac{\sigma_{LTD}^i \lambda}{\mu_{Total}}$$

Cross multiplying the constants and doing further simplification of the constants left we get the final closed form solution for the second constraint;

$$\Phi(z_i) = -\frac{c_i \mu_{total}}{C \lambda} + 1 \dots \dots \dots (36)$$

In this model too we choose the maximum value of z_i . It is obtained either by the constraint represented by equation 25 or by solving equation 36 above. The closed form solution of the model is;

$$(z_i) = \Phi^{-1} Max \left[\left(-\left[\frac{c_i \mu_{Total}}{C \lambda} \right] + 1 \right), fr_i^{-1}(MFR^i) \right]$$

APPENDIX C

ADDITIONAL GRAPHICAL RESULTS AND TABLES




Figure 36 Comparison of Type 1 service level and Fill rate for simulation of 27 products over infinite planning horizon.

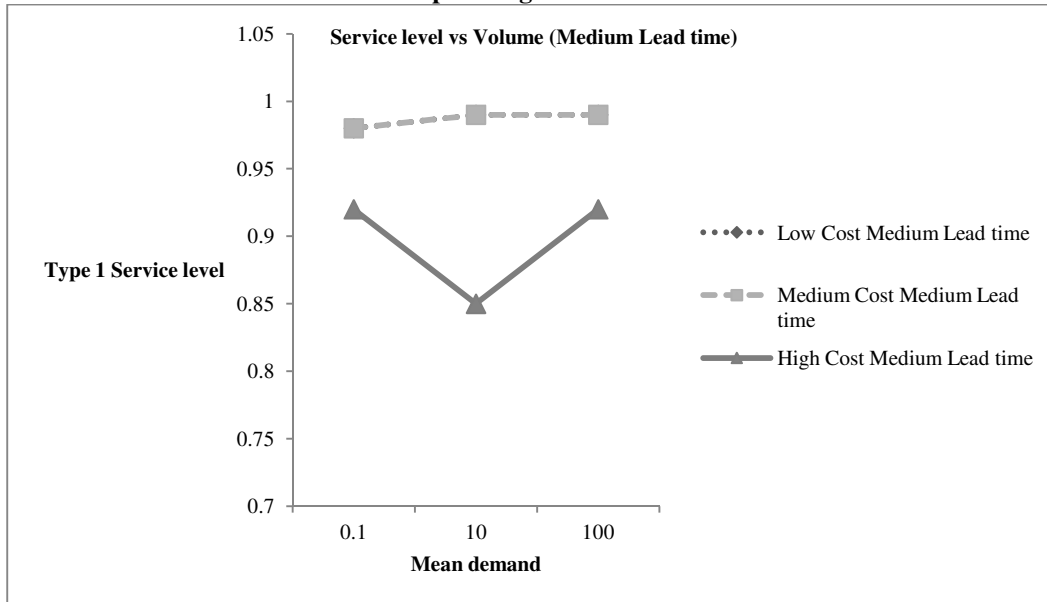


Figure 37 Change in Service level, with change in mean demand and change in cost. (Medium lead time)




Figure 38 Change in Service level, with change in mean demand and change in cost. (High lead time)




Figure 39 Change in Service level, with change in mean demand and change in lead time. (Low cost)




Figure 40 Change in Service level, with change in mean demand and change in lead time. (Medium cost)




Figure 41 Change in Service level, with change in mean demand and change in lead time. (High cost)

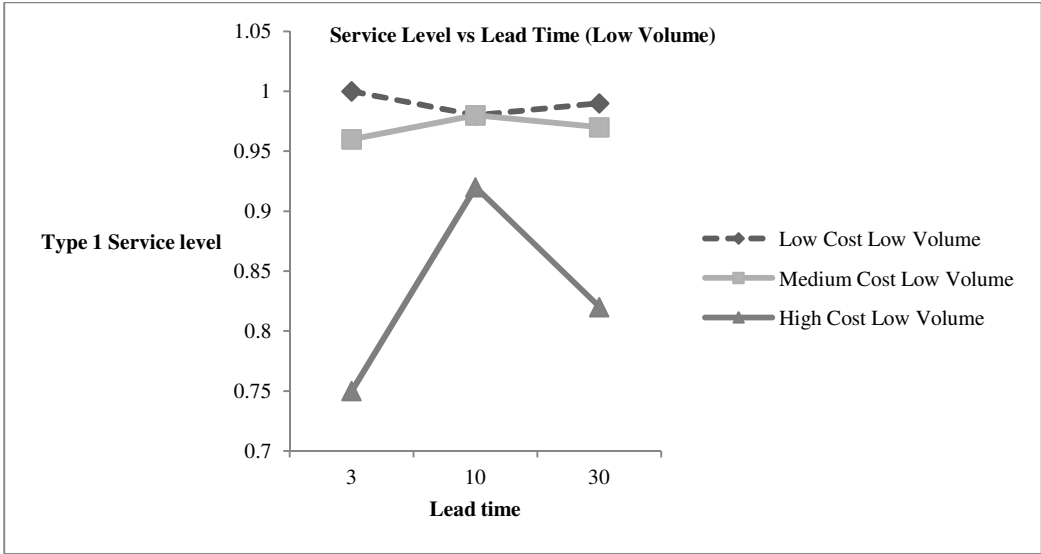


Figure 42 Change in Service level, with change in lead time and change in cost. (Low volume)




Figure 43 Change in Service level, with change in lead time and change in cost. (Medium volume)




Figure 44 Change in Service level, with change in lead time and change in cost. (High volume)




Figure 45 Change in Service level, with change in cost and change in volume. (Low Lead time)




Figure 46 Change in Service level, with change in cost and change in volume. (Medium Lead time)




Figure 47 Change in Service level, with change in cost and change in volume. (High Lead time)




Figure 48 Changes in standard deviation of Fill rate values with changes in volume of demand.




Figure 49 Changes in standard deviation of Fill rate values with changes in lead time.




Figure 50 Changes in standard deviation of Fill rate values with changes in cost.

Product No	Type 1 service level			Fill rate		
	Solution	Simulation (Poisson approximation)	Simulation (Normal approximation)	Solution	Simulation (Poisson approximation)	Simulation (Normal approximation)
10	0.99	1.00	0.99	0.99	0.97	1.00
11	0.99	1.00	0.99	0.99	0.98	0.99
12	0.85	0.84	0.79	0.93	0.69	0.91
13	0.99	1.00	0.99	0.99	0.98	0.99
14	0.99	1.00	0.99	0.99	0.94	0.99
15	0.85	0.85	0.80	0.96	0.77	0.85
16	0.99	1.00	0.99	0.99	0.99	0.99
17	0.99	1.00	0.97	0.99	0.93	0.98
18	0.85	0.84	0.82	0.98	0.78	0.98
19	0.99	1.00	0.99	0.99	1.00	1.00
20	0.99	1.00	0.97	0.99	1.00	1.00
21	0.96	0.96	0.96	0.99	0.95	0.99
22	0.99	1.00	0.99	0.99	1.00	1.00
23	0.99	1.00	0.99	0.99	1.00	1.00
24	0.92	0.96	0.91	0.99	0.93	0.98
25	0.99	1.00	0.99	0.99	1.00	1.00
26	0.99	1.00	0.99	0.99	1.00	0.99
27	0.85	0.85	0.84	0.99	0.90	0.99

Table 22 Comparison of Type 1 service level and Fill rates obtained from simulations under poisson and normal distribution with the expected values from solution for products with medium to high demand.




Figure 51 Type 1 service level variations for Product 2 for planning horizons of 1, 2 and 4 years.




Figure 52 Type 1 service level variations for Product 4 for planning horizons of 1, 2 and 4 years.




Figure 53 Type 1 service level variations for Products 5 for planning horizons of 1, 2 and 4 years.




Figure 54 Type 1 service level variations for Product 7for planning horizons of 1, 2 and 4 years.




Figure 55 Type 1 service level variations for Product 8 for planning horizons of 1, 2 and 4 years.




Figure 56 Fill rate variations for Product 1 for planning horizons of 1, 2 and 4 years.




Figure 57 Fill rate variations for Product 2 for planning horizons of 1, 2 and 4 years.




Figure 58 Fill rate variations for Product 3 for planning horizons of 1, 2 and 4 years.




Figure 59 Fill rate variations for Product 4 for planning horizons of 1, 2 and 4 years.




Figure 60 Fill rate variations for Product 5 for planning horizons of 1, 2 and 4 years.




Figure 61 Fill rate variations for Product 6 for planning horizons of 1, 2 and 4 years.

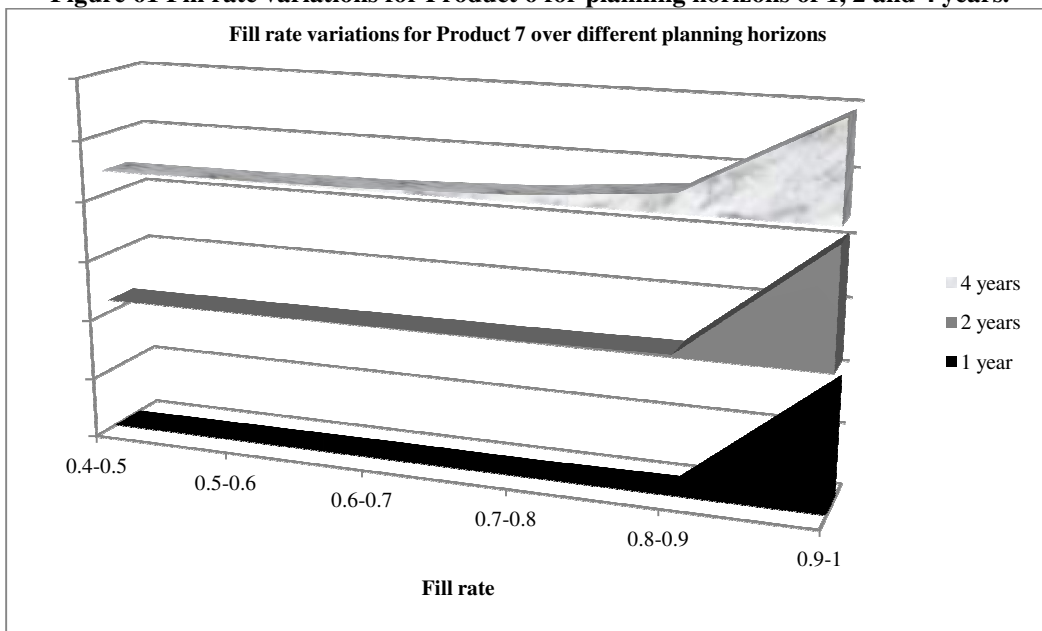


Figure 62 Fill rate variations for Product 7 for planning horizons of 1, 2 and 4 years.




Figure 63 Fill rate variations for Product 8 for planning horizons of 1, 2 and 4 years.




Figure 64 Fill rate variations for Product 9 for planning horizons of 1, 2 and 4 years.

Product No	Type 1 service level from simulation. Mean Service Level (Y^-)	Confidence interval of 95 percent. (minimum-maximum)	Minimum Type 1 service level	Maximum Type 1 service level
10	0.9993	(0.9942-1.0000)	0.9941	0.9994
11	0.9999	(0.9932-0.9999)	0.9912	0.9987
12	0.8812	(0.8801-0.8824)	0.8432	0.9976
13	0.9999	(0.9990-1.0000)	0.9845	0.9954
14	0.9937	(0.9924-0.9979)	0.9712	1.0000
15	0.9932	(0.9920-0.9987)	0.9914	0.9978
16	0.9951	(0.9940-0.9960)	0.9842	0.9998
17	0.9937	(0.9921-0.9941)	0.9845	0.9988
18	0.8471	(0.8412-0.8488)	0.8111	0.9999
19	0.9954	(0.9942-0.9962)	0.9842	1.0000
20	0.9784	(0.9723-0.9841)	0.9542	1.0000
21	0.9954	(0.9950-0.9965)	0.9874	0.9999
22	0.9944	(0.9940-0.9965)	0.9911	1.0000
23	0.9991	(0.9941-0.9951)	0.9898	0.9999
24	0.9958	(0.9950-0.9999)	0.9912	1.0000
25	0.9942	(0.9940-1.0000)	0.9823	1.0000
26	0.9991	(0.9943-1.0000)	0.9954	1.0000
27	0.8933	(0.8871-0.8981)	0.8430	0.9620

Table 23 95 percent confidence intervals, minimum and maximum Type 1 service level values for medium and fast moving products for a planning horizon of 4 years.

Product No	Fill rate from simulation. Mean Service Level (Y^-)	Confidence interval of 95 percent. (minimum-maximum)	Minimum Fill rate	Maximum Fill rate
10	0.9984	(0.9942-1.0000)	0.9842	0.9994
11	0.9931	(0.9932-0.9999)	0.9878	0.9999
12	0.9621	(0.9610-0.9705)	0.9589	0.9974
13	0.9942	(0.9980-1.0000)	0.9900	0.9954
14	0.9952	(0.9914-0.9979)	0.9938	1.0000
15	0.9990	(0.9970-0.9999)	0.9971	0.9999
16	0.9993	(0.9940-0.9999)	0.9965	0.9999
17	0.9942	(0.9921-0.9999)	0.9912	1.0000
18	0.9061	(0.8912-0.9085)	0.8989	0.9999
19	1.0000	(0.9942-1.0000)	0.9980	1.0000
20	0.9981	(0.9823-1.0000)	0.9923	1.0000

21	1.0000	(0.9950-1.0000)	0.9967	0.9999
22	0.9980	(0.9840-0.9999)	0.9974	1.0000
23	0.9998	(0.9941-0.9999)	0.9924	0.9999
24	0.9924	(0.9920-0.9999)	0.9918	1.0000
25	0.9985	(0.9940-1.0000)	0.9980	1.0000
26	0.9999	(0.9943-1.0000)	0.9980	1.0000
27	0.9732	(0.9715-0.9750)	0.9600	0.9754

Table 24 95 percent confidence intervals, minimum and maximum Fill rate values for medium and fast moving products for a planning horizon of 4 years.

REFERENCES

- Aardal, K., Jonsson, O. and Jonsson,H. 1989. Optimal inventory policies with service-level constraints. *Journal of the Operational Research Society* 40 (1):65-73.
- Altay Nezih, and Lewis Litteral. 2011. In *Service Parts Management: Demand Forecasting and Inventory Control*. First ed., Springer – Verlag London Limited.
- Alexopoloulos, C., and Seila,A.F. 1998. “Output data analysis”, Chapter 7 in *Handbook of Simulation*, J. Banks ed., Wiley, New York.
- Balci, O. 2003. Validation, verification and certification of modeling and simulation applications. *ACM Transactions on Modeling and Computer Simulation* 11 (4). 352- 377.
- Banerjee, A. and Paul,A. 2005. Average Fill rate and horizon length. *Operations Research Letters* 33 (5): 525-530.
- Bashyam, S. and Fu,M. 1998. Optimization of (s, S) inventory systems with random lead times and a service level constraint. *Management Science* 44 (12): 243-256.
- Boyaci, T. and Gallego,G. 2001. Serial production/distribution systems under service level constraints. *Manufacturing Service Operations Management* 3 (1): 43-50.
- Brooks, R. B. S., Gillen,C. A. and Lu, J.Y. 1969. *Alternative measures of supply performance: Fills, backorders, operational rate, and NORS*. Rand Corporation Memorandum, RM-60904-PR (August).

Caggiano, K. E., Jackson,P.E., Muckstadt,J.A. and Rappold,J.M. 2007. Optimizing service parts inventory in a multiechelon, multi-item supply chain with time-based customer service-level agreements. *Operations Research* 55 (2) (March 1): 303-318.

Caglar, D.,Li,C.H., and Simchi-Levi,D. 2004. Two-echelon spare parts inventory system subject to a service constraint. *IIE Transactions* 36 (7): 655-666.

Carson, J.S. (2002). Model Verification and Validation, in *Proceedings of the Winter Simulation Conference*, ed. By E. Yucesan, C-H. Chen, J.L. Snowdon, and J.M. Charnes, pp. 52-58. Association for Computing Machinery, New York.

Chen, F. Y and Krass,D. 2001. Inventory models with minimal service level constraints. *European Journal of Operational Research* 134 (1): 120-140.

Chen, J., Lin,D.K.J., and Thomas,D. 2003. On the single item Fill rate for a finite horizon. *Operations Research Letters* 31 (2): 119-123.

Cohen, M. A. and Lee,H. 1990. Out of touch with customer needs? Spare parts and after sales service. *MIT Sloan Management Review* 31 (2): 55.

Cohen, M. A., Kleindorfer,P.R., and Lee,H.L. 1989. Near-optimal service constrained stocking policies for spare parts. *Operations Research* 37 (1) (Jan. - Feb.): 104-117.

Donselaar, K.H. and Broekmeulen, R,A.C.M. 2011.Determination of safety stocks in a lost sales inventory system with periodic review, positive lead time, lot sizing and a target Fill rate. *International Journal of Production Economics*. doi:10.1016/j.ijpe.2011.05.020

Edward A. Silver, Rein Peterson. 1979. Decision systems for individual items under probabilistic demand. In *Decision systems for inventory management and production planning*. Volume- Three, 251-251, John Wiley & Sons.

Ettl, M., Feigin, G.E., Lin, G.Y. and Yao, D.D. 2000. A supply network model with base-stock control and service requirements. *Operations Research* 48 (2): 216-232.

Fu, M.C. 2002. Optimization for simulation: Theory vs practice. *Inform Journal on Computing* 14.192-215.

G. Hadley, and T.M. Whitin. 1963. *Analysis of inventory systems*. Prentice Hall.

Gianpaolo Ghiana, Gilbert Laporte, Roberto Musamanno. 2004. Introducing logistics systems. In *Introduction to logistics systems planning and control.*, ed. Richard Webber Sheldon Ross. First edition, 7-14, 15 John Wiley & Sons Ltd.

Graves, S.C. 1982. A multiple-item inventory model with a job completion criterion. *Management Science* 28 (11) (November): 1334-1337.

Hopp, Wallace J., Mark L. Spearman, and Rachel Q. Zhang. 1997. Easily implementable inventory control policies. *Operations Research* 45 (3) (May - June): 327-340.

Inderfurth, K. and Minner, S. 1998. Safety stocks in multi-stage inventory systems under different service measures. *European Journal of Operational Research* 106 (1): 57-73.

Jerry Banks, John S. Carson II, Barry L. Nelson, David M. Nicol, P. Shahabudeen. 2007. In *Discrete-Event System Simulation*. Fourth edition, Prentice-Hall: Pearson.

Johnson, M. Eric, Hau L. L., Davis, T., and Hall, R. 1995. Expressions for item Fill rates in periodic inventory systems. *Naval Research Logistics* 42 (1): 57-80.

Kleinjnen, J.P.C. 1995. Theory and Methodology: Verification and Validation of Simulation Models. *European Journal of Operational Research* 82 (1), 145-162.

Kumar, A., and Sharman, G. 1992. We love your product, but where is it? *McKinsey Quarterly*(1) (03): 24-44.

Kwon, I. H., Sung-Shick, K., and Jun-Geol, B. 2006. A simulation based heuristic for serial inventory systems under fill-rate constraints. *The International Journal of Advanced Manufacturing Technology* 31 (3): 297-304.

Larsen, C. and Thornstenson, A. 2007. A comparison between the order and the volume Fill rate for a base-stock inventory control system under a compound renewal demand process. *Journal of the Operational Research Society* 59 (6): 798-813.

Law M. Averill, and David Kelton. 2000. In *Simulation Modeling and Analysis*. Third ed., McGraw-Hill, New York.

Liu, L. and Cheung, K.L. 1997. Service constrained inventory models with random lifetimes and lead times. *Journal of the Operational Research Society* 48 (10): 1022-1028.

Mitchell, J. C. 1988. Multi-item inventory systems with a service objective. *Operations Research* 36 (5): 747-754.

Ravichandran, N. 2007. A finite horizon inventory model: An operational framework. *International Journal of Production Economics*, 108 (1-2): 406-415.

Samii, A. B., Pibernik,R., and Yadav,P. 2011. An inventory reservation problem with nesting and Fill rate-based performance measures. *International Journal of Production Economics*,133: 393-402.

Schneider, H. 1978. Methods for determining the re-order point of an (s, S) ordering policy when a service level is specified. *Journal of the Operational Research Society*: 1181-1193.

Schneider, H. and Ringuest,J.L. 1990. Power approximation for computing (s, S) policies using service level. *Management Science* 36 (7) (Jul.): 822-834.

Schrady, D. A. and Choe, U. C. 1971. Models for multi-item continuous review inventory policies subject to constraints. *Naval Research Logistics Quarterly* 18 (4): 451-463.

Schroeder, C. 1974. Managerial inventory formulations with stockout objectives and fiscal constraints. *Naval Research Logistics Quarterly* 21 (3): 375-388.

Schwarz, L. B., Deuermeyer,B.L. and Badinelli,R.D. 1985. Fill-rate optimization in a one-warehouse N-identical retailer distribution system. *Management Science* 31 (4): 488-498.

Sherbrooke, C. 1968. Metric: A multi-echelon technique for recoverable item control. *Operations Research* 16 (1) (January 1): 122-141.

Silver, E. A. 1970. A modified formula for calculating customer service under continuous inventory review. *IIE Transactions* 2 (3): 241-245.

Sobel, M. J. 2004. Fill rates of single-stage and multistage supply systems. *Manufacturing Service Operations Management* 6 (1): 41-52.

Song, J.S. 1998. On the order Fill rate in a multi-item, base-stock inventory system. *Operations Research* 46 (6) (Nov. - Dec.): 831-845.

Steven Nahmias. 2005. Inventory control subject to uncertain demand. In *Production and operations analysis*. Fifth ed., 255-261 McGraw-Hill Irwin Publication.

Thomas, D. J. 2005. Measuring item fill-rate performance in a finite horizon. *Manufacturing Service Operations Management* 7 (1): 74-80.

Thonemann, U. W., Brown, A.O., and Hausman, W.H. 2002. Easy quantification of improved spare parts inventory policies. *Management Science* 48 (9): 1213-1225.

Tijms, H. C. and H. Groenevelt. 1984. Simple approximations for the reorder point in periodic and continuous review (s, S) inventory systems with service level constraints. *European Journal of Operational Research* 17 (2) (8): 175-190.

Vliegen, I. M. H and Houtum, G.J. 2009. Approximate evaluation of order Fill rates for an inventory system of service tools. *International Journal of Production Economics* 118 (1): 339-351.

Wallace J.Hopp, Mark L. Spearman. 2001. Production scheduling. In *Factory physics*. Second ed., 488-489, McGraw Hill.

Zhang, J. and Zhang, J. 2007. Fill rate of single-stage general periodic review inventory systems. *Operations Research Letters* 35 (4): 503-509.

Zhaohui Zeng, A. and Hayya, J.C. 1999. The performance of two popular service measures on management effectiveness in inventory control. *International Journal of Production Economics* 58 (2): 147-158.