



**NAVAL
POSTGRADUATE
SCHOOL**

MONTEREY, CALIFORNIA

THESIS

**A SIMULATION OF ALTERNATIVES FOR
WHOLESALE INVENTORY REPLENISHMENT**

by

Geoffrey F. Roth

March 2016

Thesis Advisor:
Second Reader:

Arnold Buss
Javier Salmeron

Approved for public release; distribution is unlimited

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE March 2016		3. REPORT TYPE AND DATES COVERED Master's thesis
4. TITLE AND SUBTITLE A SIMULATION OF ALTERNATIVES FOR WHOLESALE INVENTORY REPLENISHMENT			5. FUNDING NUMBERS	
6. AUTHOR(S) Geoffrey F. Roth				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING /MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A			10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government. IRB Protocol number ___N/A___.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited			12b. DISTRIBUTION CODE	
13. ABSTRACT (maximum 200 words) The Navy Supply Systems Command holds \$21 billion of inventory across 430,000 repair part line items in support of ongoing naval operations. Selecting the correct reorder point for each of those line items requires a careful balance between tying up precious funding by holding too much inventory that may not be issued, and not holding enough inventory to meet customer demands in a timely fashion. This thesis compares three different methods for selecting the reorder point. The first method is a calculation that offers no attempt at optimization but is simple to understand. The second method is provided by a contractor and uses some optimization with unknown algorithmic details. The last method is a mixed-integer, linear optimization model. Comparative Inventory Simulation, a discrete event simulation model, is designed to find fill rates achieved for each National Item Identification Number under the different reorder methods. We find that average fill rates are higher in 22 out of 24 cases, and average backorder lengths are up to 50% shorter, using the last method.				
14. SUBJECT TERMS Inventory; simulation; event graphs; reorder point; fill-rate; backorder; discrete event simulation; wholesale inventory optimization model			15. NUMBER OF PAGES 89	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified		18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UU

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited

**A SIMULATION OF ALTERNATIVES FOR WHOLESALE INVENTORY
REPLENISHMENT**

Geoffrey F. Roth
Lieutenant Commander, United States Navy
B.S., University of Missouri Rolla, 2001
M.S., University of Missouri Rolla, 2003

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN OPERATIONS RESEARCH

from the

**NAVAL POSTGRADUATE SCHOOL
March 2016**

Approved by: Dr. Arnold Buss
Thesis Advisor

Dr. Javier Salmeron
Second Reader

Dr. Patricia Jacobs
Chair, Department of Operations Research

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

The Navy Supply Systems Command holds \$21 billion of inventory across 430,000 repair part line items in support of ongoing naval operations. Selecting the correct reorder point for each of those line items requires a careful balance between tying up precious funding by holding too much inventory that may not be issued, and not holding enough inventory to meet customer demands in a timely fashion. This thesis compares three different methods for selecting the reorder point. The first method is a calculation that offers no attempt at optimization, but is simple to understand. The second method is provided by a contractor, and uses some optimization with unknown algorithmic details. The last method is a mixed-integer, linear optimization model. Comparative Inventory Simulation, a discrete event simulation model, is designed to find fill rates achieved for each National Item Identification Number under the different reorder methods. We find that average fill rates are higher in 22 out of 24 cases, and average backorder lengths are up to 50% shorter, using the last method.

THIS PAGE INTENTIONALLY LEFT BLANK

THESIS DISCLAIMER

The reader is cautioned that computer programs developed in this research may not have been exercised for all cases of interest. While every effort has been made within the time available to ensure that the programs are free of computational and logic errors, they cannot be considered validated. Any application of these programs without additional verification is at the risk of the user.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

I.	INTRODUCTION.....	1
	A. BACKGROUND	1
	B. OBJECTIVES	2
	C. THESIS FLOW	3
II.	LITERATURE REVIEW	5
	A. INVENTORY POLICY.....	5
	B. DISCRETE EVENT SIMULATION	6
	C. TYPES OF SIMULATION MODELS.....	8
	D. ANALYSIS OF SIMULATION MODELS	9
	E. WHOLESALE INVENTORY OPTIMIZATION MODEL	10
	1. Optimization	10
	2. The WIOM Model.....	11
III.	COMPARATIVE INVENTORY SIMULATION DEVELOPMENT AND DATA SETS.....	15
	A. ASSUMPTIONS.....	15
	B. CIS FORMULATION	15
	1. Diagram and Description of Events	15
	2. Inputs to CIS	18
	3. Outputs from the Model.....	19
	C. DATA SETS	19
IV.	ANALYSIS	23
	A. EFFECTS OF WARMUP ON THE RESULTS	23
	B. COMPARISON OF FILL-RATES ACROSS DATA SETS.....	29
	1. Overall Results	30
	2. Fill-rates Achieved by NIIN	33
	3. Selection of Demand Probability Distributions.....	36
	C. AVERAGE TIME BACKORDERED	37
	D. DISTRIBUTION ANALYSIS.....	41
	E. SENSITIVITY TO -TIME.....	47
V.	CONCLUSIONS AND RECOMMENDATIONS FOR FOLLOW-ON STUDIES	49
	A. CONCLUSIONS	49
	B. RECOMMENDATIONS FOR FOLLOW-ON RESEARCH.....	50

APPENDIX A. SUPPLY TERMS51

APPENDIX B. FILL-RATE GRAPHS BY DATASET53

APPENDIX C. FILL-RATE DISTRIBUTION GRAPHS59

LIST OF REFERENCES63

INITIAL DISTRIBUTION LIST65

LIST OF FIGURES

Figure 1.	Example of an Order-point, Order-quantity (s, Q) System.....	6
Figure 2.	Simple Event Graph.....	7
Figure 3.	Example of an Event Graph.....	8
Figure 4.	The Three Steps to Calculate the DPDF	12
Figure 5.	Resulting DPDF	12
Figure 6.	CIS Event Graph.....	18
Figure 7.	S5S8 Fill-rate for One Million Days, No Warmup.....	24
Figure 8.	S5S8 Fill-rate for 1.4 Million Days, 400,000 Day Warmup.....	25
Figure 9.	Non-nuclear Consumables: 100 Most Variable, No Warmup	26
Figure 10.	Non-nuclear Consumables: 100 Most Variable, Warmup.....	27
Figure 11.	Non-nuclear Consumables: 100 Highest Variance to Mean, No Warmup.....	28
Figure 12.	Non-nuclear Consumables: 100 Highest Variance to Mean, Warmup.....	29
Figure 13.	Aviation Consumables - Fill-rates Achieved by NIIN (Empirical)	34
Figure 14.	Aviation Repairables - Fill-rates Achieved by NIIN (Empirical).....	35
Figure 15.	Aviation Repairables - WIOM Fill-rate Difference for All Distributions.....	36
Figure 16.	Aviation Repairables - WIOM Fill-rate Difference by Distribution	37
Figure 17.	Backorder Length - Empirical	39
Figure 18.	Backorder Length - Parametric	40
Figure 19.	Backorder Example.....	41
Figure 20.	Difference Between Estimated and Simulated Fill-rates	42
Figure 21.	Five-Factor Classification Tree for Aviation Repairables	45
Figure 22.	Two Factor Classification Tree for Aviation Repairables	46
Figure 23.	Maritime Consumables: Fill-rates Achieved by NIIN (Empirical)	53
Figure 24.	Maritime Repairables: Fill-rates Achieved by NIIN (Empirical).....	54
Figure 25.	Aviation Consumables: Fill-rates Achieved by NIIN (Parametric).....	55
Figure 26.	Aviation Repairables: Fill-rates Achieved by NIIN (Parametric)	56
Figure 27.	Maritime Consumables: Fill-rates Achieved by NIIN (Parametric).....	57
Figure 28.	Maritime Repairables: Fill-rates Achieved by NIIN (Parametric)	58

Figure 29.	Aviation Repairables: WIOM Fill-rate Difference for All Distributions.....	59
Figure 30.	Aviation Repairables: WIOM Fill-rate Difference by Distribution.....	59
Figure 31.	Maritime Consumables: WIOM Fill-rate for All Distributions.....	60
Figure 32.	Maritime Consumables: WIOM Fill-rate by Distribution	60
Figure 33.	Maritime Repairables: WIOM Fill-rate Difference for All Distributions.....	61
Figure 34.	Maritime Repairables: WIOM Fill-rate Difference by Distribution.....	61

LIST OF TABLES

Table 1.	Probability Distributions and Estimators to Approximate Lead-time Demand.....	13
Table 2.	Description of Data Sets Provided by NAVSUP.....	20
Table 3.	Sample Fill-rate Calculation.....	31
Table 4.	Fill-rate Results Using Empirical Demand.....	32
Table 5.	Fill-rate Results Using Parametric Fit.....	33
Table 6.	Backorder Length – Empirical.....	39
Table 7.	Backorder Length – Parametric.....	40
Table 8.	Fill-rate Over-estimation and Under-estimation Across All Datasets.....	43
Table 9.	Abbreviation Used in Classification Trees.....	44
Table 10.	Classification Groupings.....	44
Table 11.	Five Factor Classification Tree Results.....	45
Table 12.	Two Factor Classification Tree Results.....	46
Table 13.	Sensitivity of the Fill-rate to Stochastic Lead-time.....	47

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS AND ABBREVIATIONS

CIS	Comparative Inventory Simulation
DPDF	Demand probability distribution function
LSSI	Level Setting Strategy Indicators
MIP	Mixed integer program
NAVSUP	Navy Supply Systems Command
NIIN	National Item Identification Number
PD	Probability Distribution
SC	Simple Calculation
SPO	Service Planning and Optimization
WIOM	Wholesale Inventory Optimization Model

THIS PAGE INTENTIONALLY LEFT BLANK

EXECUTIVE SUMMARY

The Navy Supply Systems Command (NAVSUP) holds \$21 billion of inventory across 430,000 repair part line items in support of ongoing naval operations. In an unconstrained world, NAVSUP would hold enough inventory to support repair of all naval systems through their entire life cycle; unfortunately, budget and storage constraints force NAVSUP to carry a limited number of repair parts and consumable items known as NIINs (National Item Identification Numbers) in its inventory. This process involves the selection, for each NIIN, of an inventory position quantity at which to reorder the NIIN, known as the “reorder point.” Determination of optimal reorder points for all NIINs considered is a critical factor in supporting the warfighter and maintaining material readiness.

This thesis compares three methods for selecting reorder points. A first method suggested by NAVSUP is the use of a simple calculation (SC) based on mean and standard deviation of lead-time demand. The second method is a contractor-provided tool called Service Planning and Optimization (SPO), currently in use at NAVSUP, which algorithmic details are unknown. The last method is the Wholesale Inventory Optimization Model (WIOM), a mixed-integer, linear optimization program recently developed at the Naval Postgraduate School. This thesis explores the effectiveness of the three different methods for calculating reorder points by developing a discrete-event simulation tool called Comparative Inventory Simulation (CIS).

CIS is designed to model wholesale demands, issues, and reorders. CIS can be broken down into ten different events (such as order placing and receive, demand, etc.). As events occur, they modify the value of variables in the system (such as inventory levels, backorders, etc.) and schedule future events. Time proceeds in variable increments to the time of the next event. CIS uses a number of inputs from each NIIN, the most important two being the reorder point (as determined by one of the three methods) and the demand probability distribution (PD) function describing the number of demand events that occur during a time period. This PD is generated either from a parametric fit (based on mean and variance of the demand), or from an empirical distribution (based on

historical monthly demand). Using the PD, CIS schedules events such as filling demands, and placing and receiving reorders. In this process, CIS counts the number of backorders for each NIIN, which allows the calculation of fill rates. At the conclusion of the simulation, CIS reports the fill rate and the mean length of time backordered for each NIIN.

Five different datasets supplied by NAVSUP are explored via simulation: S5S8, aviation consumables, aviation repairables, maritime consumables, and maritime repairables. Each dataset contains a variety of high- and low-demand items as well as different PDs representing the demand patterns. The number of NIINs in each dataset ranges from 249 to 9,601.

CIS can use an initial “warmup” period in order to determine the time required to overcome initialization biases (due to initial inventories). We conduct analysis of replications with and without using the warmup period, and find that in both cases fill rates achieve steady state (remain nearly constant within a 1% band) around the 400,000 day in the simulation.

The primary method of evaluating the effectiveness of the three methods of reorder point selection has been by comparing the fill rates achieved by simulating the demands for each NIIN. For each dataset, the overall, average fill rate and weighted fill rate are calculated using both empirical and parametric-fit PDs resulting in 24 total cases. Simulated fill-rates using WIOM’s reorder points lead to the highest observed fill rates in 22 out of the 24 cases. The fill rates resulting from the use of the SC reorder points were only higher than SPO in a single case.

A second consideration is that WIOM’s ability to estimate the fill rates is mixed and largely dependent upon the PD used for each NIIN. WIOM estimates the fill-rates within 2% of the simulated value for approximately 25% of the NIINs, and CIS indicates that WIOM is twice as likely to over-estimate as to under-estimate the fill-rate. When the demand PDs are fit to a Poisson distribution, WIOM is generally accurate. Fits made with the gamma PD are consistently less accurate.

CIS has uncovered significant differences in the mean length of backorders between reorder points selected by WIOM and SPO. Simulated lengths are up to 50% shorter when using WIOM. It appears SPO places the reorder points for some NIINs too high at the cost of not being able to buy sufficient safety stock for other NIINs. WIOM spreads the safety stock more evenly throughout the NIINs. As a result, WIOM has fewer lengthy backorders.

THIS PAGE INTENTIONALLY LEFT BLANK

ACKNOWLEDGMENTS

The research contained in this thesis could not have been accomplished without the direction, patience, and expertise of Professor Arnold Buss and Professor Javier Salmeron. Their guidance during the simulation development and analysis as well as their support during the writing of the paper were invaluable.

I would also like to thank my family. My wife, Sarah, provided encouragement to write one more page every night. Nathan provided comic relief. Abby was the example of boundless energy and enthusiasm. Without all of you, this would not have been possible.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

A. BACKGROUND

The Navy Supply Systems Command (NAVSUP) holds \$21 billion of inventory across 430,000 repair part line items in support of ongoing naval operations. The Chief of Naval Operations has directed all NAVSUP activities to establish sparing requirements in order to meet readiness objectives (Office of the Chief of Naval Operations 2011). Finding the correct amount of spares to hold for each line item under uncertain demand and financial constraints is a critical factor in determining how effectively NAVSUP can support the warfighter (Silver 1998). In an unconstrained world, NAVSUP would hold enough inventory to support repair of all naval systems through their entire life cycle; unfortunately, budget and storage constraints force NAVSUP to carry a diminished number of repair parts in its inventory and select an inventory quantity at which to reorder the item. Determination of the best reorder point for each repair part is a critical factor in supporting the warfighter and maintaining material readiness.

NAVSUP considers fill rate to be the primary metric to measure the effectiveness of the inventory policy at the wholesale level. Fill rate is defined using the following equation:

$$\text{fill rate} = \frac{\text{total demands} - \text{number backordered}}{\text{total demands}}$$

This formula can be applied to a specific National Item Identification Number (NIIN) or to a group of NIINs. Each NIIN is assigned a fill-rate target based on the relative criticality of the item. For instance, NAVSUP may assign a circuit card used to repair a radar a 90% fill-rate target while a plastic gauge cover may be assigned a 70% fill-rate target. By assigning different fill-rate targets to different NIINs, NAVSUP is able to prioritize those items that must be available for issue above those items for which some amount of delay is acceptable.

The most straightforward method for calculating the reorder point uses the simple calculation (SC) as follows (Oswald 2014):

$$\text{reorder point} = \text{mean demand over lead-time} + t\sqrt{\text{variance of the lead-time demand}}$$

In SC, “t” is a value selected to make the reorder points small enough such that the total sum of the value of the safety stock inventory met the budget constraint. The value of “t” is adjusted upward or downward as appropriate until the budget constraint is achieved. As neither fill rate nor any other metric are considered in the selection of “t,” the reorder points, while easy to calculate, are clearly not optimal.

NAVSUP currently uses Morris Cohen Associates Service Planning and Optimization (SPO) model to determine the number of each line item to carry in its inventory. As a contractor-provided solution, SPO is largely a “black box,” and the details of how the optimum inventory levels are determined are unknown. Additionally, anecdotal observations suggest that SPO favors buying inexpensive items to maximize fill rate while minimizing costs. By exceeding fill-rate goals with inexpensive items, SPO is able to maintain a high overall average fill rate while falling short on the fill-rate goals of expensive items. While this may make perfect sense in a business setting where a corporation seeks to maximize profit, it is preferable to meet all fill rate goals (as allowed by the constraints) in a military readiness environment. As a result, individual item managers must manually adjust reorder point for some NIINs further degrading the optimality of the solution.

In an effort to provide a lower-cost alternative to SPO, Salmeron and Craparo (2014) developed the Wholesale Inventory Optimization Model (WIOM). WIOM simultaneously selects reorder values across all NIINs with the goal of minimizing weighted deviations from fill-rate targets. It uses historic data about each NIIN as input into a General Algebraic Modeling Language implementation of a mixed-integer programming (MIP) model and produces a reorder point for each NIIN as output. Additional information about WIOM is located in Chapter II.

B. OBJECTIVES

This research develops Comparative Inventory Simulation (CIS), a discrete-event simulation of demands and replenishments in order to answer the following questions:

- Are the fill-rate estimation calculations performed by WIOM accurately reflected by the simulated fill-rates?
- How do the reorder points selected by SC, SPO and WIOM affect the overall fill-rate?
- Does the selection of the reorder point by SPO or WIOM change the amount of time it takes for a backordered NIIN to be filled?
- Does the selection of a demand probability distribution affect fill-rate deviation?

This thesis provides qualitative and quantitative information to help NAVSUP explore the effects of using different methods to determine the reorder point. The result of this thesis will assist NAVSUP in determining the strengths and weaknesses of the SC, SPO, and the more recently developed WIOM, and inform future decision making with regard to different reorder point strategies.

C. THESIS FLOW

Chapter II contains the literature review and explores work that has already been completed on similar problems. Chapter III describes the assumptions used in developing CIS, CIS itself and the datasets used to explore the problems. Chapter IV contains the results of the experiments. Chapter V draws conclusions from the results and provides some recommendations for follow-on work.

THIS PAGE INTENTIONALLY LEFT BLANK

II. LITERATURE REVIEW

It is helpful to conduct a review of previous studies and publications to establish a foundation upon which the simulation can be constructed and that data can be analyzed. First, we examine different inventory policies that could be used to support the Navy's goals. We will look at how discrete-event simulation could be applied to the problem. We conclude the literature review with a survey of the work concerning the analysis of simulation models. The definition of inventory terms used in this section are located in Appendix A.

A. INVENTORY POLICY

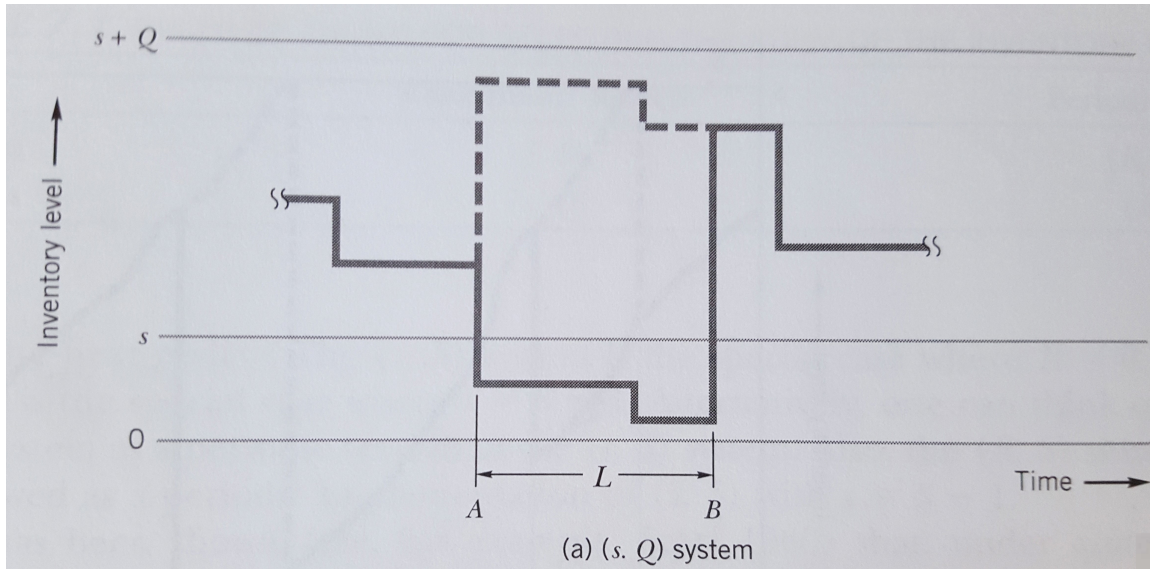
One of the fundamental problems faced by inventory managers is when a replenishment order should be placed for an item in their inventory (Silver et al. 1998). Placing an order too soon results in excess inventory and ties up capital that could be used to hold other items. Placing an order too late causes inventory shortfalls and subsequent delays in providing the end-user with the material they require. To address this problem, NAVSUP utilizes an order-point, order-quantity (s, Q) system.

The order-point, order-quantity inventory policy is a continuous review system where the inventory position for a particular NIIN has the ability to trigger a reorder event. In this system, inventory position is defined as the quantity of a particular item held in inventory plus the quantity currently on order minus the backordered quantity. As items are issued from inventory, inventory position decreases. When inventory position is equal to the reorder point (s), a reorder event for quantity Q occurs (see Figure 1).

Silver et al. (1998) notes several advantages achieved when the order-point, order-quantity system is used. Since inventory position considers ordered but unreceived items, this inventory policy precludes placing an extra order today when material arriving tomorrow would be sufficient to cover the shortage. Additionally, the system can be easily visualized. The primary drawback to the order-point, order-quantity system is that there is not a defined procedure to handle those cases when individual demand quantities exceed the reorder quantity such that the reorder of Q items will not raise inventory

position above the reorder point. For this thesis, orders are assumed to always be for a single quantity so the drawback is not a problem.

Figure 1. Example of an Order-point, Order-quantity (s, Q) System



Source: Silver EA, Pkye DF, Peterson R. (1998). *Inventory Management and Production Planning and Scheduling*. (John Wiley & Sons, New York). The solid line in this figure represents on-hand inventory, and the dotted line represents inventory position. Items are issued from the inventory until the inventory position drops below the reorder point (s) at time A when an order in the amount of the reorder quantity (Q) is placed. At time B , the outstanding order is received. Since there are not any outstanding orders, inventory position is equal to on-hand inventory.


B. DISCRETE EVENT SIMULATION

Discrete event simulation utilizes state changes that occur at a discrete set of points along a time axis (Schruben 1983, Law 2007, Sanchez 2007). As events occur, they modify the value of variables that describe the system. Events may also schedule future events to occur. Events are held in an event list as a data pair while waiting execution. The first part of the pair is the name of the scheduled event. The second part of the pair is the time at which the event will occur. This list is maintained in a time-sorted queue, and at the conclusion of the current event, the next event is popped off the queue and simulation time is advanced to the time of next scheduled event. Thus, the simulation proceeds from one event to another. Since state transitions only occur at events, the time

between events can be passed over without concern. Note that simulated time does not proceed in fixed but variable time increments that are determined by when events are scheduled.


The relationship between events and the changes that they make on the system can be described by event graphs. Event graphs consist of two different elements: event nodes and scheduling edges. Each event node corresponds with a state transition that transpires in the system. During an event, state variables are changed according to the state transition function for that event. Event nodes are connected by scheduling edges. These edges can be marked with conditional statements and time delays. If the scheduling edge has a conditional statement, the connected event node will only be scheduled if the conditional statement is true. Conditional statements are represented on the event graph as an equation or inequality in parentheses. If the scheduling edge has a time delay, the connected event node will be scheduled to occur after the time delay has elapsed. Time delays are written as the letter t with a subscript. See Figures 2 and 3.

Figure 2. Simple Event Graph



Source: Schruben, Lee. 1983. Simulation Modeling with Event Graphs. In *Communications of the ACM* 26 (11): 957–963. This event graph indicates that t time units after the occurrence of event j , event k will be scheduled to occur providing condition (i) holds at the time event j occurs.

Figure 3. Example of an Event Graph



In this event graph, arrival are scheduled every t_a time units. When an arrival occurs, a start service event is scheduled if Q is less than one. If Q is one or greater the arrival is discarded. Upon the start service event, Q is increased by one and the end service event is scheduled t_s time units later. When the end service event occurs, Q is decreased by one. This event graph represents a system where a person arrives for service. If they can immediately be helped, service begins. If they would have to wait, they leave.

C. TYPES OF SIMULATION MODELS

There are two different types of simulation models that could be considered for inventory analysis: terminating and nonterminating. A terminating simulation is one for which there is a natural time horizon over which the simulation should be run. The simulation ends when one of three criteria are met: the system is cleared out, the point at which no further useful information is obtained, or a time specified by management mandate (Law 2007). A system is cleared out is when the last event occurs and the queue is empty. Suppose a store has set hours of operation. The simulation would start when the store opens, and the simulation would terminate when the store closes. An example of a simulation that terminates without further information would be a simulation of a race. Data would be collected up to the point when the last racer crossed the finish line, at which point additional data collection would be meaningless. An example of a simulation that terminates on a time specified by mandate is a simulation, which is to collect data for seven days. After the seventh day, the simulation ends.

A nonterminating simulation is one for which there is no event to specify the length of a run (Law 2007). This is most useful when there is concern about the long-term behavior of a system but there no specific definition of long-term exists. In these

simulations, the desire is to find the steady state of the variables of the measures of performance. Since this thesis is concerned with finding the fill-rate achieved by the selection of different reorder points, a nonterminating simulation is the natural selection for implementation.

D. ANALYSIS OF SIMULATION MODELS

There are several considerations to be taken into account when conducting analysis of a simulation model. The first problem is the sensitivity of the simulation to the starting conditions. At the beginning of the simulation, a steady state has not yet been achieved; therefore, the data collected has been influenced by whatever values have been initially assigned to the variables. Law (2007) described this phenomenon as the problem of the initial transient and proposed suspending data collection until a warmup period has passed. The problem then became determining the length of the warmup period. Welch (1983) proposed a solution to determining the length of the warmup period through a four-step process. In step one, multiple replications of the simulation were generated. In step two, the i^{th} observation from each replication were averaged to find the mean value for that observation. In step three, a weighted moving average was computed from the averaged values to smooth the high frequency oscillations. In step four, the weighted moving average was plotted and the point at which convergence has appeared is selected as the warmup period.

The second consideration concerns the effect of the input parameters on the output of the simulation. If there were a single input parameter into the model, that parameter could be varied to gauge its effect. However, if multiple input variables need to be tested, other means need to be employed. A 2^k factorial design allows for two different levels to be selected for k inputs. The simulation is then run for each combination of factor levels for each input parameter. There are two problems with this approach. The first is the restriction to two levels per input parameter, many simulations would find this too limiting. The second is the number of simulation runs required to cover all 2^k cases with enough iterations to construct confidence intervals. The computation time is simply unacceptable.

An m^k factorial design removes the two-level restriction posed by the 2^k level design (Sanchez 2008). This design allows for m levels to be selected for each input parameter thereby greatly increasing the solution space that can be searched. However, this increases the number of cases that have to be considered and is even more computationally expensive than the 2^k factor design.

Nearly Orthogonal Latin Hypercubes offer an alternative approach giving similar solution space exploring qualities as the m^k factorial design with a significantly diminished number of samples. Cioppa and Lucas (2007) proposed a method for constructing Nearly Orthogonal Latin Hypercubes.

E. WHOLESALE INVENTORY OPTIMIZATION MODEL

WIOM is the result of an effort to provide an alternative method to NAVSUP in order to find optimal reorder points for their wholesale inventory. This section will explore how WIOM uses optimization to determine these values.

1. Optimization

There are many instances where a person seeks to maximize (or minimize) a particular result. “Optimization models represent problem choices as decision variables and seek values that maximize or minimize objective functions of the decision variables subject to constraints on variable values expressing the limits on possible decision choices” (Rardin 1998). The objective function measures the effectiveness of the values of the decision variables. The value of the decision variables are limited by a set of equalities and inequalities called the constraints. The set of values for the decision variables which maximize (or minimize) the value of the objective function while satisfying the constraints represents the optimal solution.

There are three categories of optimization defined by the characteristics of the objective function and the constraints. A linear program is one in which all of the decision variables are continuous and both the constraints and the objective function are linear functions of the decision variables. A MIP has the characteristics of a linear program but relaxes the requirement for continuity by allowing the decision variables to

be restricted to integer values. A nonlinear program allows for nonlinear functions to appear in the objective function as well as the constraints.

2. The WIOM Model

According to Salmeron and Craparo (2014), “WIOM seeks reorder points that are globally optimal for all NIINs simultaneously considered, while observing a number of constraints.” This optimization model was implemented as a MIP. The MIP seeks to minimize the sum of the penalties assessed when the NIIN’s target fill rate is missed. This necessitates finding an estimate for fill rate, which, in turn, requires an estimate for the demand probability distribution function (DPDF).


One way to generate the DPDF is to leverage the empirical distribution as represented by 60 months of demand history for each NIIN. This distribution is unique for each NIIN and each value represents the number of demands for a single quantity in a particular month. A three step process converts the empirical distribution to DPDF. In step one, the monthly demand is uniformly distributed during the month. In step two, the months are converted to lead-times. Additionally, the lead-time demand \hat{x} and the cycle length \hat{c} are calculated. In step three, the lead-time demand is converted to cycle demand and the occurrences of each demand level are counted to find the DPDF. An example of this process is in Figure 4 and the resulting distribution is in Figure 5.

Figure 4. The Three Steps to Calculate the DPDF

Monthly Demand:

1	2	3	0	2	0	1	1
---	---	---	---	---	---	---	---

 Lead-time: 2 months
 Q: 1

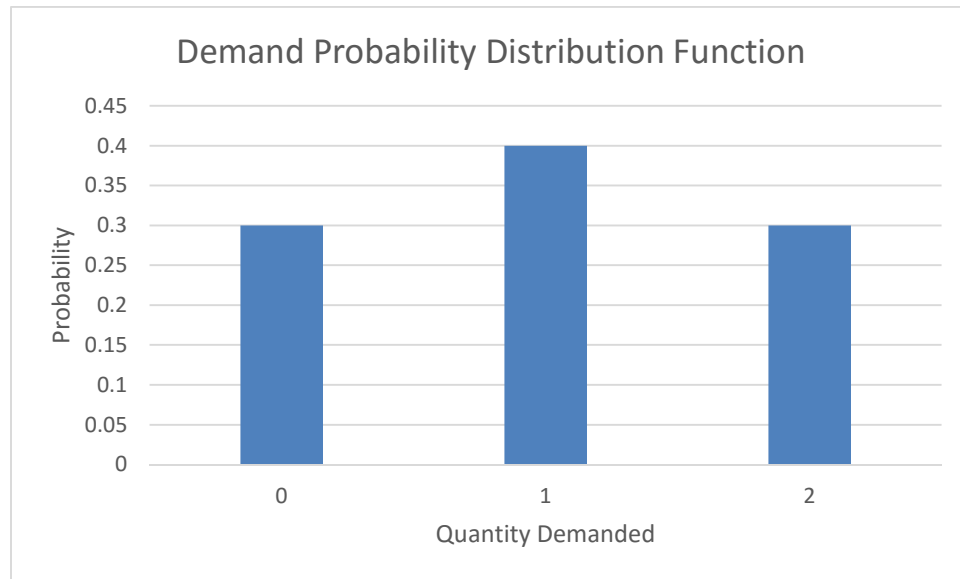


This example uses monthly demand data with a two-month lead-time. An ‘x’ in the image represents where a demand occurs in time. In the first step, the monthly demand is uniformly distributed during the corresponding month. In step two, the months are converted into lead-times. Since the length of a lead-time is two months, the eight months are converted into four lead-times. We then calculate the lead-time demand \hat{x} and the

cycle length \hat{c} .
$$\hat{x} = \frac{10 \text{ demands}}{4 \text{ lead-times}} = 2.5 \frac{\text{demands}}{\text{lead-time}} \quad \hat{c} = \max\left\{1, \frac{\hat{x}}{Q}\right\} = \max\left\{1, \frac{2.5}{1}\right\} = 2.5 \frac{\text{cycles}}{\text{lead-time}}$$

The occurrences at each demand level are counted to find the DPDF in Figure 5.

Figure 5. Resulting DPDF



This graph captures the cyclic demand from Figure 4. Out of the 10 cycles, three experienced zero demands, four experienced four demands and three experienced two demands.

Ideally, WIOM would always be able to use an empirical distribution representing the demand during the cycle time. However, three reasons prevent the empirical distribution to be used for some NIINs. First, monthly demand data are not available for all NIINs. Secondly, some NIINs have so little demand that the empirical distribution cannot be applied. Lastly, some NIINs had fractional values for demands during the month. For these cases, we use the expected demand \hat{x} and standard deviation $\hat{\sigma}$ during the lead-time. WIOM converts the lead-time parameters into cycle-time parameters during a preprocessing step. WIOM then converts these cycle-time parameters into a DPDF by applying a parametric fit to a prefixed standard distribution suggested by NAVSUP. The details of these fits can be found in the WIOM documentation, with a summary provided by Table 1.

Table 1. Probability Distributions and Estimators to Approximate Lead-time Demand

Prob. Dist.	Name (Parameters)	Mean	Variance	Variance/Mean	Parameter Estimation (Based on $\hat{x}, \hat{\sigma}$)
1	Poisson (λ)	λ	λ	1	$\lambda = \hat{x}$ or $\lambda = \hat{\sigma}^2$
2	Negative Binomial (r, p) ($r = \#$ failures to stop) (Generalized Neg. Bin.)	$rp / (1-p)$	$rp / (1-p)^2$	$1 / (1-p) > 1$	$\tilde{p} = 1 - \frac{\hat{x}}{\hat{\sigma}^2}; \tilde{r} = \frac{\hat{x}^2}{\hat{\sigma}^2 - \hat{x}}$
3	Normal (μ, σ)	μ	σ^2	σ^2 / μ	$\mu = \hat{x}; \sigma = \sqrt{\hat{\sigma}^2}$
4	Binomial (n, p)	np	$np(1-p)$	$1-p < 1$	$\tilde{p} = 1 - \frac{\hat{\sigma}^2}{\hat{x}}; \tilde{n} = \frac{\hat{x}^2}{\hat{x} - \hat{\sigma}^2}$
5	Gamma (k, θ) (k =shape, θ =scale)	$k\theta$	$k\theta^2$	θ	$k = \frac{\hat{x}^2}{\hat{\sigma}^2}; \theta = \frac{\hat{\sigma}^2}{\hat{x}}$
6	Deterministic (μ)	μ	0	0	$\mu = \hat{x}$

Source: Salmeron J, Craparo E (2014) Wholesale Inventory Optimization Models. Working paper, Naval Postgraduate School.

The estimated fill-rate calculation used in WIOM was inspired by Silver et al. (1998). However, the data provided by NAVSUP violated two of the assumptions of Silver's calculation (Salmeron et al. 2014). The first was that the average level of backorders could be significant when compared with the average level of on-hand stock.

The second was that multiple cycles could occur during a single lead-time. These two issues can be corrected by modifying Silver's calculation by substituting cycle-time demand for lead-time demand. The resulting equation allows for WIOM to estimate fill-rate during optimization.

III. COMPARATIVE INVENTORY SIMULATION DEVELOPMENT AND DATA SETS

The CIS is intended to determine the fill-rate achieved for a NIIN given a particular cycle-time demand probability distribution function, reorder quantity and reorder point. The Python computer language is used to implement the discrete event simulation.

A. ASSUMPTIONS

In order to develop CIS and test the datasets, several assumptions are required:

- The probability distribution functions determined by WIOM accurately reflect the actual demand distributions.
- When continuous distributions are used to determine the quantity demanded in a time period, the resulting random variate can be rounded.
- Lead-times for each NIIN are known and constant.
- Demand for each NIIN is independent of demands for all other NIINs.
- Individual demands are for one unit of the NIIN.
- All demands have equal priority and will be filled in the order received.

B. CIS FORMULATION

1. Diagram and Description of Events

The simulation can be represented in an event graph containing ten different events. The following notation (for a given NIIN i) is helpful in understanding the event graph and a definition of these terms is in Appendix A:

- t_i : the length of a lead-time (days)
- t'_i : the length of a cycle (days)
- s_i : the reorder point (units of issue)
- Q_i : the reorder quantity (units of issue)
- L_i : the on-hand inventory quantity (units of issue)

- a_i : the number of randomly determined demands in a particular cycle (units of issue)
- BB_i : the number of items currently backordered (units of issue)
- OO_i : the number of items currently on order (units of issue)

The event graph is shown in Figure 6 and a description of each event follows:

Run: The simulation begins with the run event scheduled. Upon execution, the run event sets the starting on-hand inventory level equal to the sum of the reorder point and the reorder quantity. This starting quantity represents the maximum amount of inventory that could be held at one time (a reorder has been received during a time when there were not any demands). While any amount greater than the reorder point could have been selected, this quantity was convenient and will later be shown not to have an impact on the results.

Cycle: At the completion of the run event, the cycle event is scheduled. In a traditional inventory model, the cycle is defined by the length of the lead-time (Silver et al. 1998). However, this definition is founded upon the assumption that the reorder quantity is greater than the lead-time demand. This assumption does not hold for the NAVSUP data. Therefore, in this simulation the length of a cycle for NIIN i is defined by

$$t'_i = \frac{t_i}{\max\left(1, \frac{\hat{x}_i}{Q_i}\right)}$$

where t_i is the lead-time, \hat{x}_i is lead-time demand and Q_i is the reorder

size. The cycle event first determines the number of demand events that will transpire during the cycle according to the probability distribution function related to the NIIN. It then schedules both the “schedule cycle” and the “schedule demand” event.

Schedule Cycle: The schedule cycle event occurs t' time units after the corresponding cycle event. This event checks to see if there is still time remaining to be simulated. If so, it schedules an additional cycle. If there is not, no further events are scheduled and the simulation will terminate.

Schedule Demand: A demand event is immediately scheduled at a random time between 0 and t' time units in the future using a uniform random distribution. The event

then decrements the number of events remaining to be scheduled by one. If there are any demands remaining to be scheduled, another schedule demand event is scheduled. For cases when multiple demands are to be scheduled, the demands could be scheduled out of order (e.g., in a 30 day cycle, the random draw for the first demand could be five days into the cycle and the random draw for the second demand could be two days into the cycle). To ensure correct execution, the event list is sorted in increasing order by time at the conclusion of this event.

Demand: The demand event corresponds to a single item demand. No variables are changed in the demand event. This event exists only as a transition point to the backorder or issue event. The demand event is required because the times from the schedule demand state could have been randomly generated out of time sequence. If on-hand inventory is greater than zero, the “issue” event is scheduled. Otherwise, the “backorder” event is scheduled.

Issue: In the issue event, the on-hand balance is decremented by one and the “order for stock” event is scheduled if the current on-hand balance is less than or equal to the reorder point.

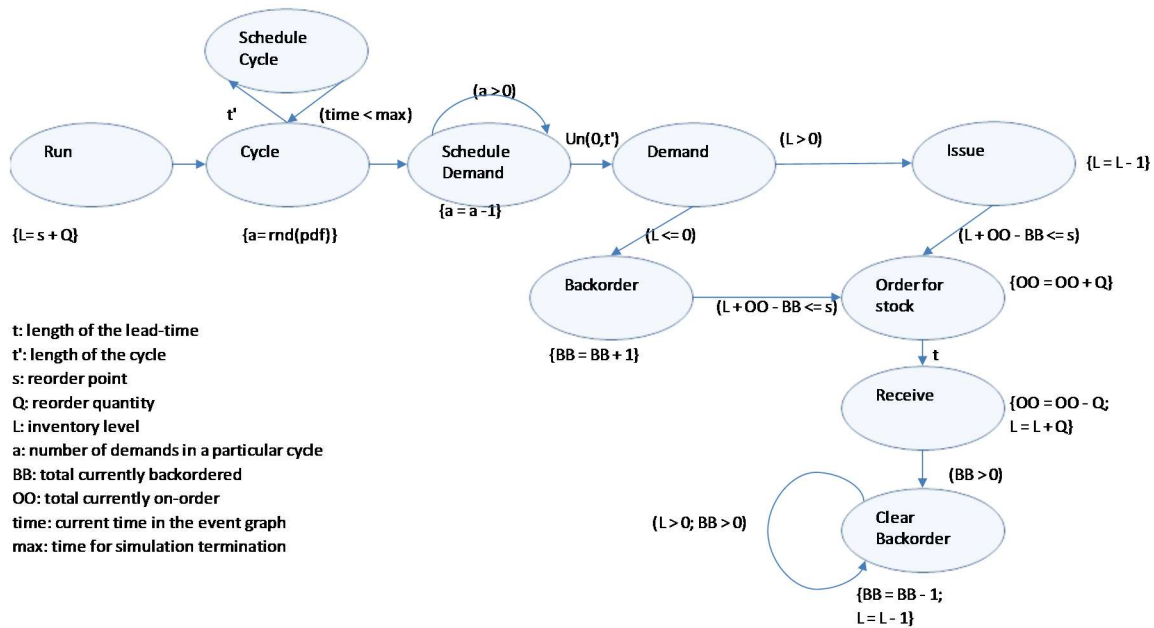
Backorder: The backorder event causes the number of items currently backordered to increment by one. The event then checks to see if an additional order needs to be placed. If the current on-hand balance plus the number on-order minus the number backordered is less than or equal to the reorder point the “order for stock” event is scheduled.

Order for stock: The order for stock event increments the number of items on order by the reorder quantity and schedules the “receive” event.

Receive: The receive event decreases the number of items on order by the reorder quantity and increases the on-hand balance by the reorder quantity. The event then checks to see if there are any items backordered; if so, the “clear backorder” event is scheduled.

Clear backorder: The clear backorder event decreases the number of items backordered and the on-hand inventory by one. If on-hand inventory and the number backordered are greater than one, the “clear backorder” event is scheduled.

Figure 6. CIS Event Graph



2. Inputs to CIS

There are a number of parameters that are input into CIS, which are used to determine the behavior of the NIIN. The reorder point defines the inventory level at which a reorder will be triggered. The reorder quantity defines that number of units of a NIIN that will be acquired in a single reorder. The cycle length defines the amount of time to which the demand probability distribution function refers. Lastly the DPDF specifies the number of demands that could occur during a cycle. There are three possible ways that the DPDF can be defined:

- **Probability distribution function** - The simulation can accept different probability distribution functions as input along with the associated parameters. The simulation supports the Poisson, negative binomial, normal, binomial, gamma, Pareto, Weibull, uniform and deterministic distributions.
- **Empirical distribution** - The empirical distribution is input as any number of consecutive months with the number of demands per month. The simulation generates a uniform random integer variate from 1 to the number of months in the empirical distribution. This random variate is used to select which month's demand is sampled for the cycle.

- Probability distribution file - While solving the for the optimal reorder points, WIOM generates a file listing probabilities for each demand quantity over one cycle..

The analysis conducted as part of this research will explore the effect of varying different values of reorder quantity, reorder point, and lead-time to explore the sensitivity of these parameters.

3. Outputs from the model

The simulation determines several important statistics during execution. The total number of demand events and the total number of times that a demand could not be filled by on-hand inventory are recorded and used to determine the fill-rate according to the formula in Chapter I. The simulation also records the means and the standard deviations of monthly and cycle demand to provide confidence that the simulation is operating properly. The minimum and maximum inventory position observed are recorded to keep track of extremes during execution. Lastly, statistics concerning the shortest, longest, mean and standard deviation of backorder times are recorded as an indication of the effectiveness of the reorder point in addition to fill-rate.

C. DATA SETS

Five different data files from January 2016 were provided by NAVSUP for exploration using the simulation. These data files contain both 60 months of empirical monthly demand data as well as the mean and standard deviation of lead time demand for the existence of the NIIN. Although not an exhaustive list of all NIINs stocked by NAVSUP, these data sets contain a mixture of high-demand and low-demand items supporting a variety of platforms. These data sets provided the reorder point determined by SPO as well as the inputs required to calculate reorder points using both WIOM and the SC. Table 2 summaries the characteristics of each data set.

Table 2. Description of Data Sets Provided by NAVSUP

	Total NIINs	Budget	Normal	Binomial	Poisson	Neg. Binomial	Gamma	Deterministic
S558	249	\$58,642	12	33	17	88	99	0
Non-nuclear Consumable	2231	\$17,189,771	150	542	140	847	552	0
Non-nuclear Repairable	7200	\$170,019,778	1	2251	714	3174	1060	0
Aviation Consumable	1516	\$7,923,000	155	239	90	527	505	0
Aviation Repairable	9601	\$887,486,000	703	2030	74	2575	4210	9

This table describes the distribution of NIINs according to the rules of the parametric fit described in Chapter II. The budget for each category is for the value of safety stock allowed for one year.

To calculate the reorder points using WIOM, the data sets were input into the WIOM Excel graphical user interface. The budget solve options were set to solve using MIP with a 1% optimality gap or a two hour runtime limit. Separate WIOM runs were conducted to find the reorder points using the empirical data files and using the parametric estimations. The resulting output files from WIOM were used as input into the simulation.

The reorder points for the SC were determined by using the aforementioned formula.

$$\text{reorder point} = \text{mean demand over lead-time} + t\sqrt{\text{variance of the lead-time demand}}$$

To find the solution for t , a Python script was created. The script takes four inputs from the user: the name of a NAVSUP provided wholesale data file as input for the NIIN's price, lead-time demand and variance; the name of a WIOM_NIIN file for modification; the output filename; and the budget constraint. Upper and lower bounds for t are set at 100 and -100, respectively. The script then selects the midpoint between the upper and lower bounds and uses that t to find the safety stock. If the safety stock is fractional, the safety stock is rounded to the nearest integer. If the safety stock less than zero, it is raised to zero. The total yearly cost of safety stock for all NIINs is found by multiplying the safety stock level of each NIIN by the price NIIN and then summing for all NIINs. If the cost exceeds the budget, the upper bound is lowered to the current t . If the cost is less than the budget, the lower bound is set to the current t . New values for t are selected until the upper and lower bounds converge to 0.01.

A possible source of error exists in the data, which may contribute to erroneous results. The error concerns a mismatch between the average demand calculated using the empirical distribution and the average demand provided as a parameter in the datasets. Since WIOM uses these parameters as inputs into the reorder point calculation, this introduces some uncertainty about the optimality of the answer provided by the model. As SPO is a black box algorithm, we are unsure if or how this error effects SPO's calculation. For the purpose of this thesis, we assume that the error affects both methods equally and a comparison between the two can still be made.

THIS PAGE INTENTIONALLY LEFT BLANK

IV. ANALYSIS

This chapter provides the significant results from the simulation study. In the first section, the effects of different length warmup times are examined to determine the appropriate simulation length for further experimentation. In the second section, fill-rates are compared across different data sets using the different reorder point picking methods. The third section examines the mean length of time a NIIN is backordered resulting from the different methods. Section four looks at WIOM's ability to estimate fill-rates for different demand distributions. In the last section, the effect of varying the lead-time is examined. This analysis uses the January 2016 input files provided by NAVSUP. The results are highly dependent upon the input values and may not be representative of the results using different inputs.

A. EFFECTS OF WARMUP ON THE RESULTS

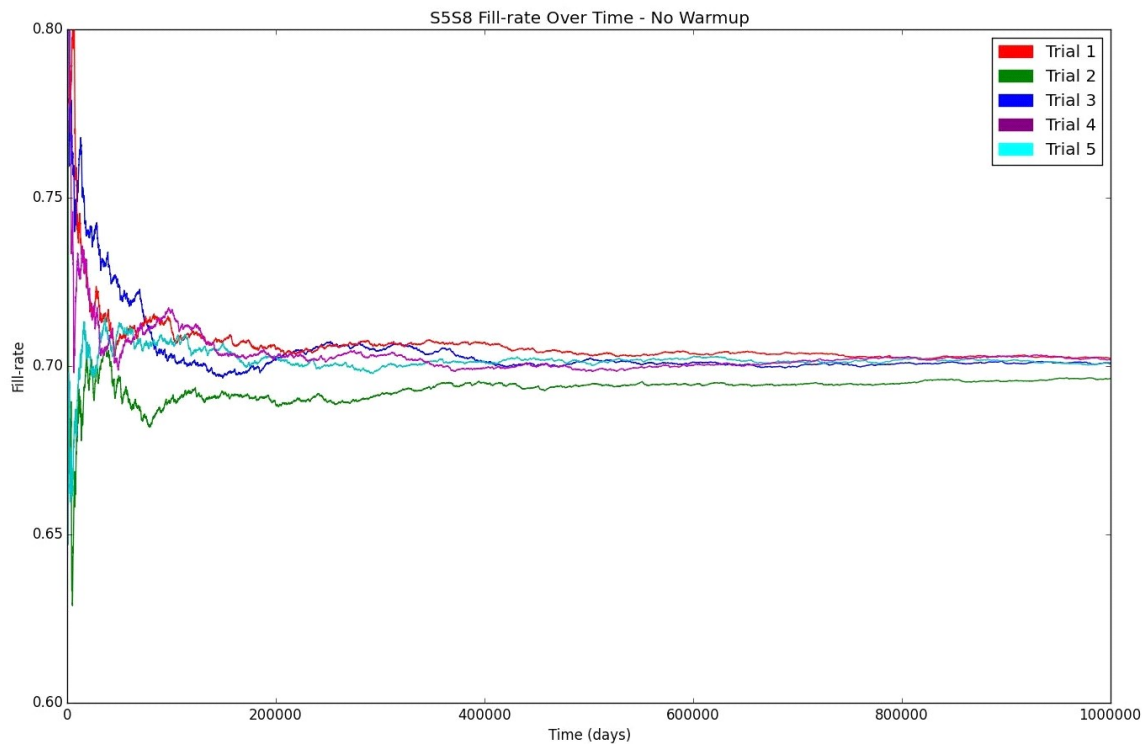
Three different trials are used to explore the effects of warmup on the simulation results. Following Welch's (1983) procedure, five different runs are conducted on each dataset. Each time an event occurs, the fill-rate is calculated and recorded. Plots of each fill-rate are constructed and the convergence point across the trials is observed. This convergence point represents the time when the simulation reached steady state and warmup effects are no longer affecting the fill-rate.

The initial five replications were conducted using the S5S8 dataset. The S5S8 dataset was selected for two reasons. The first reason was that all of the different demand probability distributions were represented. The second reason was that it was found that the amount of data recording required to plot the fill-rate graphs was such that using one of the larger datasets was computationally prohibitive. Demands were simulated for each NIIN for one million days and the fill rate over all NIINs was recorded following each event in the simulation.

As shown in Figure 7, the fill-rates for the replications reach steady state around day 400,000 in the simulation. Therefore, the simulation was repeated using a 400,000-day warmup period. Figure 8 shows that the simulation reaches steady state more quickly

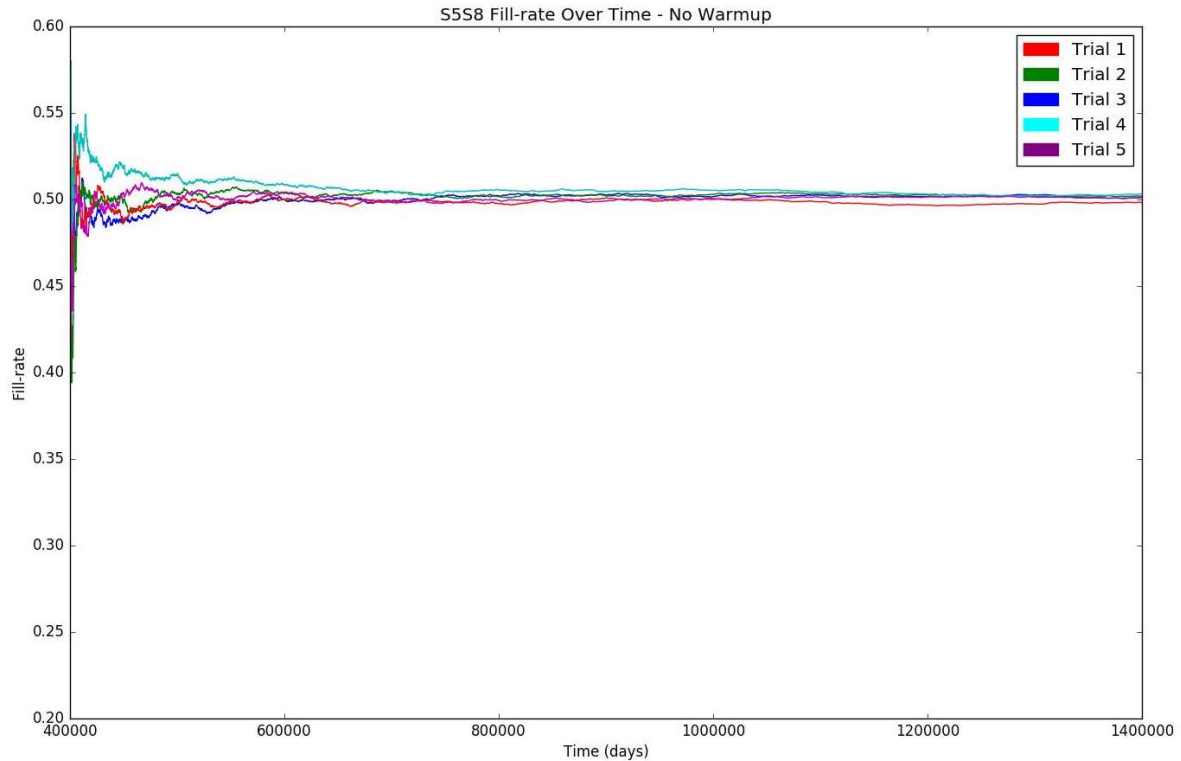
with the warmup period; however, the overall fill-rate is not affected. Even with the warmup period, the fill-rate is not in steady state immediately at the beginning of the simulation. This is an expected behavior. Although starting biases have been removed, the fill-rate calculation is a running average of binary events, and a demand is either met via an issue from stock or the demand is backordered. It will take some number of events to occur for the fill-rate to reach steady state.

Figure 7. S5S8 Fill-rate for One Million Days, No Warmup



Five replications of run length one million days are run using the S5S8 data set. The trials converge to steady state about day 400,000.

Figure 8. S5S8 Fill-rate for 1.4 Million Days, 400,000 Day Warmup



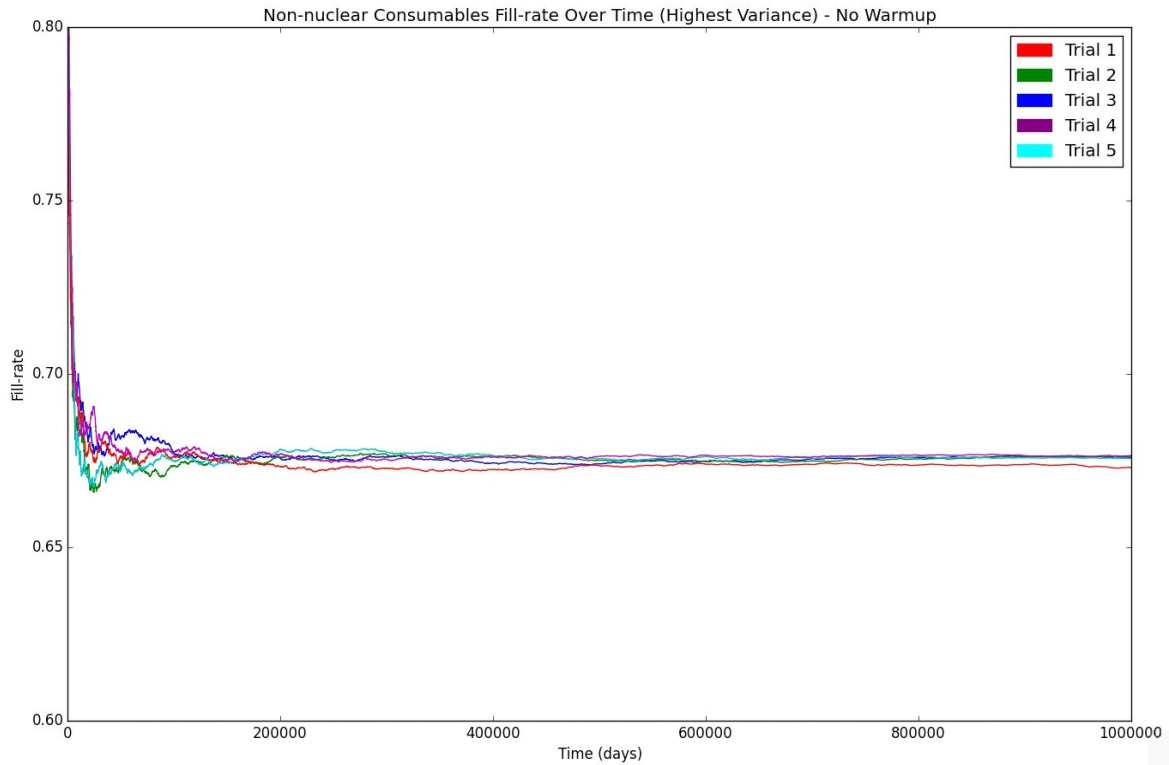
Five replications with run length 1,400,000 days are run using the S5S8 data set. A warmup period of 400,000 days is specified. In this case, the trials converged more quickly than without the warmup period; however, the steady state fill rate observed did not change.

These set of results suggested two things. The first was that one million days was sufficiently long in order to observe a reasonable accurate measurement of the fill-rate of the system. The second was that a warmup period was not necessary in this simulation with this number of days simulated to find a good estimate of the fill-rate. To further solidify these claims, two additional cases were tested.

The next case used the one hundred NIINs from the non-nuclear consumable dataset with the highest variance. Those NIINs with a high variance may take longer to converge than those NIINs with a smaller variance. Again, two sets of trials each with five replications were conducted. The first trial was one million days long without a warmup period. As with the previous dataset, Figure 9 shows that fill-rate reached steady state in about 400,000 days. In Figure 10, using a 400,000 day warmup period to remove

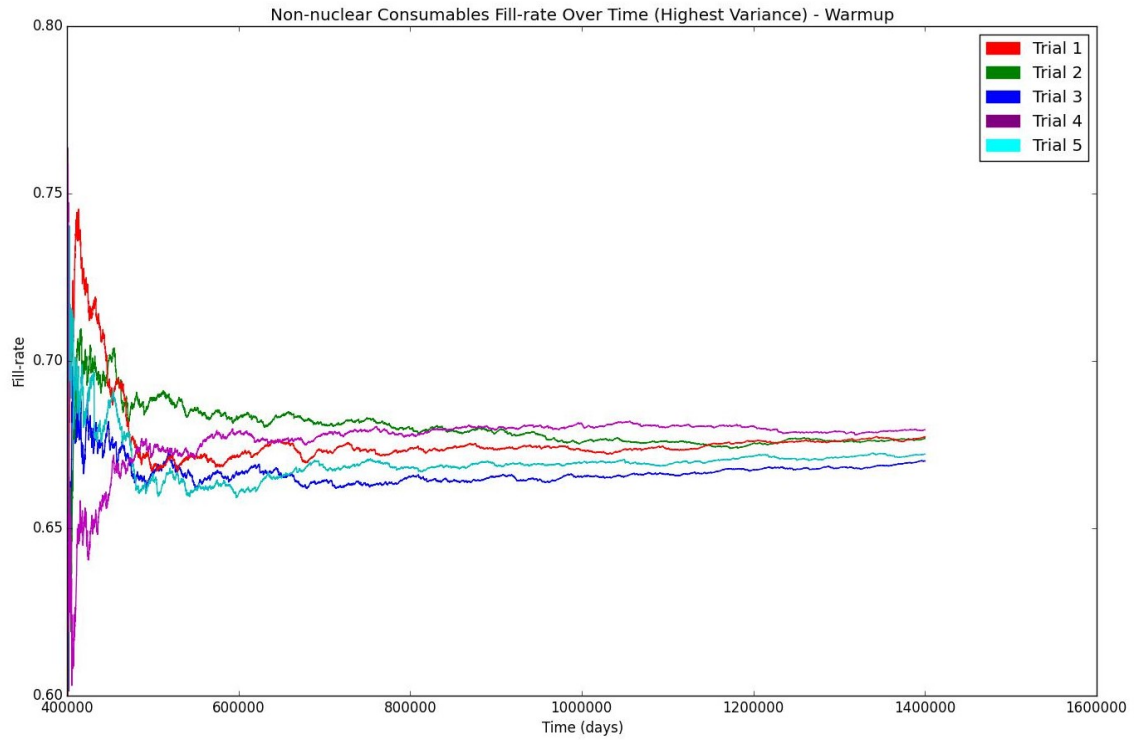
starting biases did not have an effect on the final fill-rates. In order to be completely confident that warmup periods can be neglected, one more set of experiments were run.

Figure 9. Non-nuclear Consumables: 100 Most Variable, No Warmup



Five replications are conducted using the 100 NIINs with the highest variance from the non-nuclear consumable data set. The fill rate reaches steady state at 400,000 days without a warmup period.

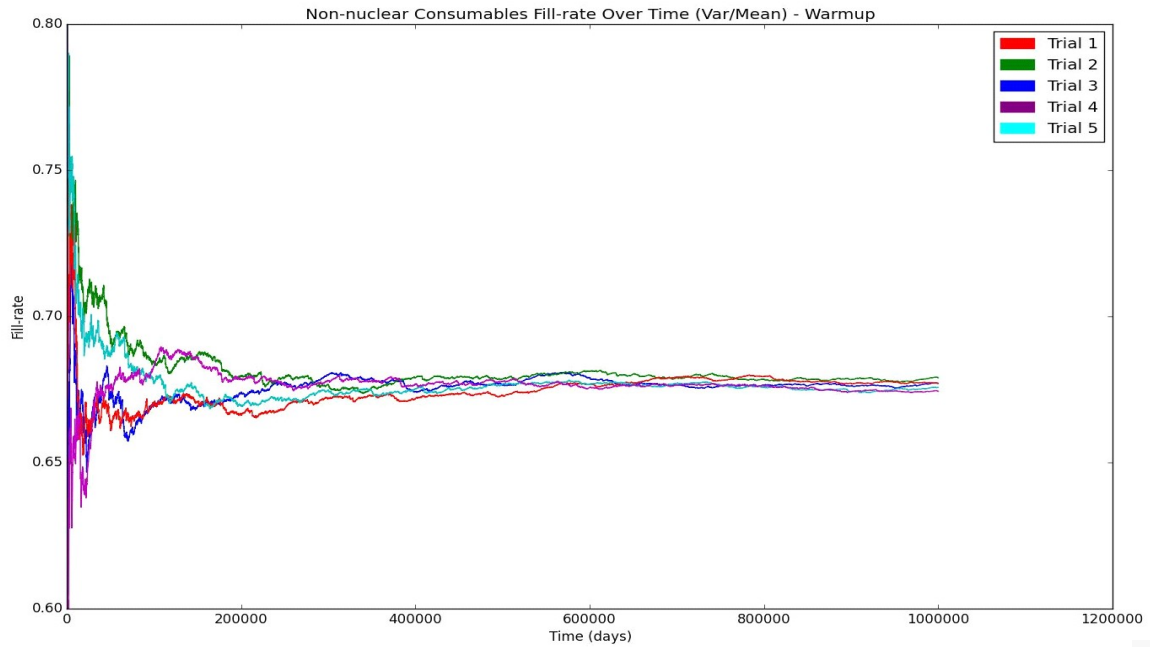
Figure 10. Non-nuclear Consumables: 100 Most Variable, Warmup



Five replications are conducted using the non-nuclear consumable data set with a 400,000 day warmup period. The fill rates reach steady state after 300,000 days of observations and the steady state fill rate value has not changed from the no-warmup period case.

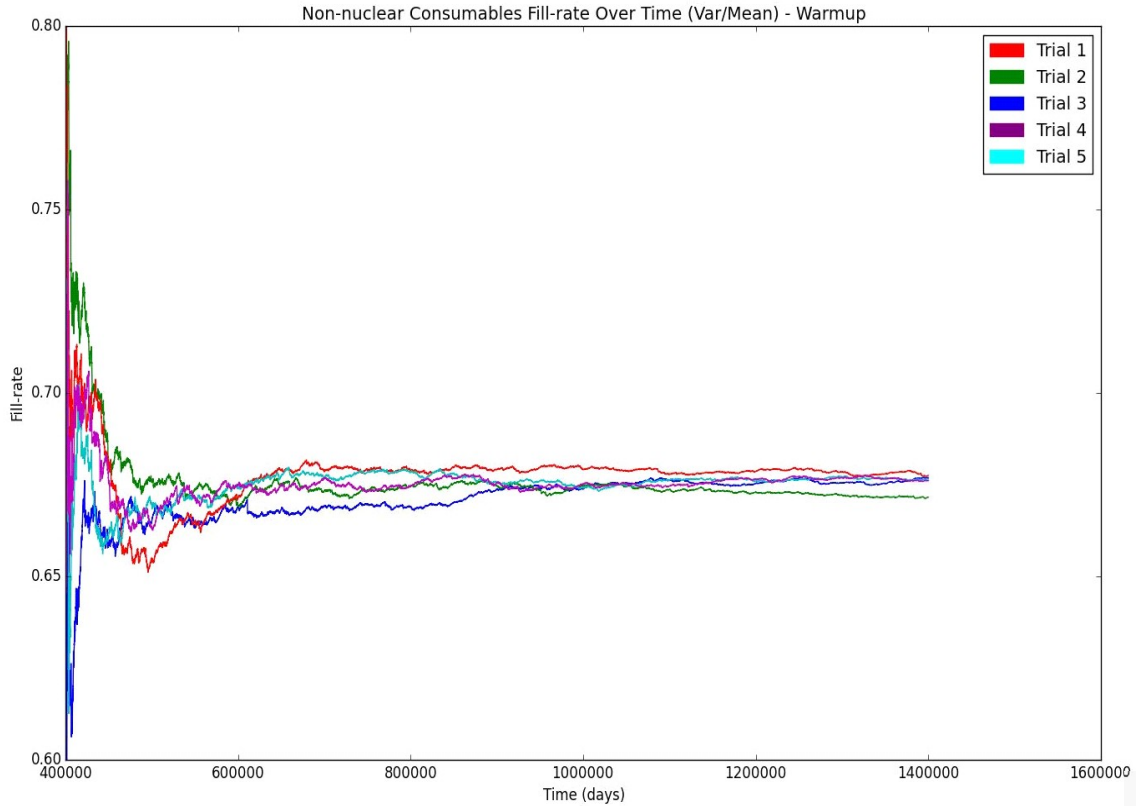
In a modification of the second set of trials, the last set selected the one hundred NIINs with the largest variance to mean ratios. Once again, the non-nuclear consumable dataset was selected. Such ratios for this subset ranged from 2.5 to 14.2.

Figure 11. Non-nuclear Consumables: 100 Highest Variance to Mean, No Warmup



Five replications were conducted using 100 NIINs with the highest variance to mean ratios from the non-nuclear consumable data set. As we observed before, steady state was achieved after 400,000 days.

Figure 12. Non-nuclear Consumables: 100 Highest Variance to Mean, Warmup



Five replications are conducted using a 400,000-day warmup period with 100 NIINs with the highest variance to mean ratios from the non-nuclear consumable data set. The steady state value of fill rate has not changed from the no-warmup case.

Since a warmup period did not seem to affect the fill-rates observed in any of the three datasets, a warmup period was not used for future experimentation.

B. COMPARISON OF FILL-RATES ACROSS DATA SETS

The most important question this thesis seeks to explore concerns the effectiveness of the three different methods (WIOM, SPO and SC) in achieving high fill-rates. To explore this question, four different datasets were explored: aviation consumables, aviation repairables, maritime consumables and maritime repairables. The methods were applied to each dataset in order to find the reorder point for each NIIN. Ten replications of the simulation are conducted for each method's reorder points using both the empirical and parametric methods for scheduling demands. In general, a larger

number of replications would be preferred. However, even taking advantage of parallel processing, each additional 10 replications would take seven days of processing time to accomplish. Time limitations did not allow for these additional runs to occur.

The usual problems that occur from a limited number of runs were offset by choosing run lengths that were long enough that the simulation reached steady state. This allowed the simulation to overcome the initial transient problem. Additionally, as the simulation is in steady state, we would not expect significant changes in the value of the output. The observed standard deviations between the 10 runs on each dataset were very small and there is no reason to believe additional runs would have any effect on the overall results.

1. Overall Results

A general formula for calculating fill rate is presented in Chapter 1. Here we show three different refinements to the fill rate calculation: overall fill rate, average fill rate, and weighted fill rate. Overall fill rate is defined as follows for NIINs i :

$$\text{fill rate}_{\text{overall}} = \frac{\sum_i \text{demands}_i - \sum_i \text{backordered}_i}{\sum_i \text{demands}_i}$$

Additionally, the average fill-rate for NIINs i is calculated as:

$$\text{fill rate}_{\text{average}} = \frac{\sum_i \text{fill rate}_i}{\text{total number of NIINs}}$$

Finally, weighted fill rate for NIINs i is:

$$\text{fill rate}_{\text{weighted}} = \frac{\sum_i \text{fill rate}_i * \frac{\text{lead time demand}_i}{\text{lead time}_i}}{\sum_i \frac{\text{lead time demand}_i}{\text{lead time}_i}}$$

Overall fill rate and average fill rate are affected by outliers but are sensitive in different ways. In the average fill-rate calculation, low demand items have an outsized effect on the value since all NIINs are weighted equally regardless of demand. The overall fill-rate

calculation essentially weights the NIINs relative to their demand but allows the effects of low demand NIINs to be dominated by the results of the high demand NIINs. An example of these effects can be seen in Table 3. The weighted fill rate is a similar to the overall fill rate in that it takes the magnitude of the number of items demanded into account. The difference between the two calculation is that the weights in weighted fill rate are calculated using the mean lead-time demand where the overall fill rate uses the observed demand in the simulation. The results of all three calculations are presented to provide an overall picture of how well the methods satisfy demand. It is important to note that WIOM only estimates weighted fill rate.

Table 3. Sample Fill-rate Calculation

NIIN _{<i>i</i>}	Demand _{<i>i</i>}	Backordered _{<i>i</i>}	Fill Rate _{<i>i</i>}
1	1	1	0.00%
2	10	1	90.00%
3	20	3	85.00%
4	1000	0	1.00%

overall fill-rate: 99.61%

average fill-rate: 44.00%

This table shows the difference between the overall fill-rate and the average-fill rate. The overall fill-rate is heavily influenced by the high demand of item 4. In contrast, the average fill-rate is heavily influenced by the single failure to fill of NIIN 1.

Tables 4 and 5 summarize the simulated fill rates found using the empirical demand distributions and the parametric fit, respectively. In these tables, the mean and the standard deviation represent the values found across the 10 replications, not the differences between NIINs in a single replication.

Table 4. Fill-rate Results Using Empirical Demand

		Empirical Distribution							
		Aviation Consumables		Aviation Repairables		Maritime Consumable		Maritime Repairable	
Overall Fill		Mean	Standard Dev.	Mean	Standard Dev.	Mean	Standard Dev.	Mean	Standard Dev.
WIOM		0.72479	0.00035	0.72893	0.00016	0.67207	0.00021	0.73194	0.00032
SPO		0.33623	0.00115	0.64419	0.00020	0.68903	0.00051	0.51633	0.00034
SC		0.34462	0.00088	0.58166	0.00000	0.59064	0.00026	0.41186	0.00001
Weighted Fill									
WIOM Projected		0.89000		0.91000		0.81000		0.81000	
WIOM		0.72849	0.00026	0.72902	0.00017	0.66962	0.00021	0.80669	0.00018
SPO		0.34413	0.00038	0.64422	0.00024	0.68789	0.00041	0.51774	0.00051
SC		0.34829	0.00029	0.60493	0.00056	0.58770	0.00069	0.41332	0.00028
Average Fill									
WIOM		0.94328	0.00013	0.85334	0.00012	0.73832	0.00027	0.80682	0.00012
SPO		0.50808	0.00037	0.65112	0.00014	0.57650	0.00005	0.46873	0.00012
SC		0.50057	0.00042	0.64345	0.00025	0.51627	0.00015	0.45026	0.00002

Four datasets are used to test the fill-rate results using the empirical demand distribution. In one of the four cases, the weighted fill-rate observed is nearly identical to the WIOM estimate. In ten out of the twelve cases, significantly higher fill-rates were simulated using the reorder points selected by WIOM that those selected by SPO or SC.

The first thing to notice about the results under the empirical demand distribution is how small the standard deviations are for the outcome of the 10 runs. This adds additional weight to the suggestion that the run length of the simulation is appropriate to reach steady state. In the maritime repairable case, WIOM was able to estimate the average fill-rate achieved through simulation with only a 1% deviation. However, in the other three cases, WIOM over-estimated the fill rate up to a 19% difference.

In all WIOM cases, the average fill-rate is higher than the overall fill-rate. This indicates the WIOM is doing a better job selecting reorder points for low demand items (weighted more heavily in the average fill-rate calculation) than for high demand items (weighted more heavily in the overall fill-rate calculation). In contrast, the fill-rates achieved by SPO are higher (in two of four cases) for the overall fill-rate suggesting that SPO is favoring higher demand items at the expense of filling low demand orders. It is unclear why the fill rates achieved by SPO are so poor for the aviation consumable data.

Table 5. Fill-rate Results Using Parametric Fit

		Parametric Fit						
		Aviation Repairables		Maritime Consumable		Maritime Repairable		
Overall Fill	on Consur	Mean	Standard Dev.	Mean	Standard Dev.	Mean	Standard Dev.	
WIOM		0.87746	0.74557	0.00012	0.75913	0.00069	0.75267	0.00035
SPO		0.32281	0.64629	0.00006	0.68731	0.00009	0.51581	0.00009
SC		0.30025	0.64483	0.00008	0.58739	0.00005	0.40854	0.00008
Weighted Fill								
WIOM Projected		0.92000	0.91000		0.86000		0.88000	
WIOM		0.87705	0.74716	0.00046	0.75993	0.00021	0.75337	0.00017
SPO		0.32277	0.65092	0.00032	0.68599	0.00007	0.57170	0.00018
SC		0.30041	0.64493	0.00009	0.65273	0.00014	0.40858	0.00011
Average Fill								
WIOM Projected		0.95793	0.92770		0.85054		0.84677	
WIOM		0.94348	0.84166	0.00046	0.79086	0.00021	0.80948	0.00017
SPO		0.50787	0.65137	0.00032	0.57923	0.00007	0.46916	0.00018
SC		0.49906	0.68362	0.00009	0.51444	0.00014	0.45011	0.00011

Four data sets are used to find the fill rates resulting from the use of the three different methods to select reorder points. The reorder points selected by WIOM resulted in higher fill rates than the reorder points select by the other two methods in twelve out of twelve cases.

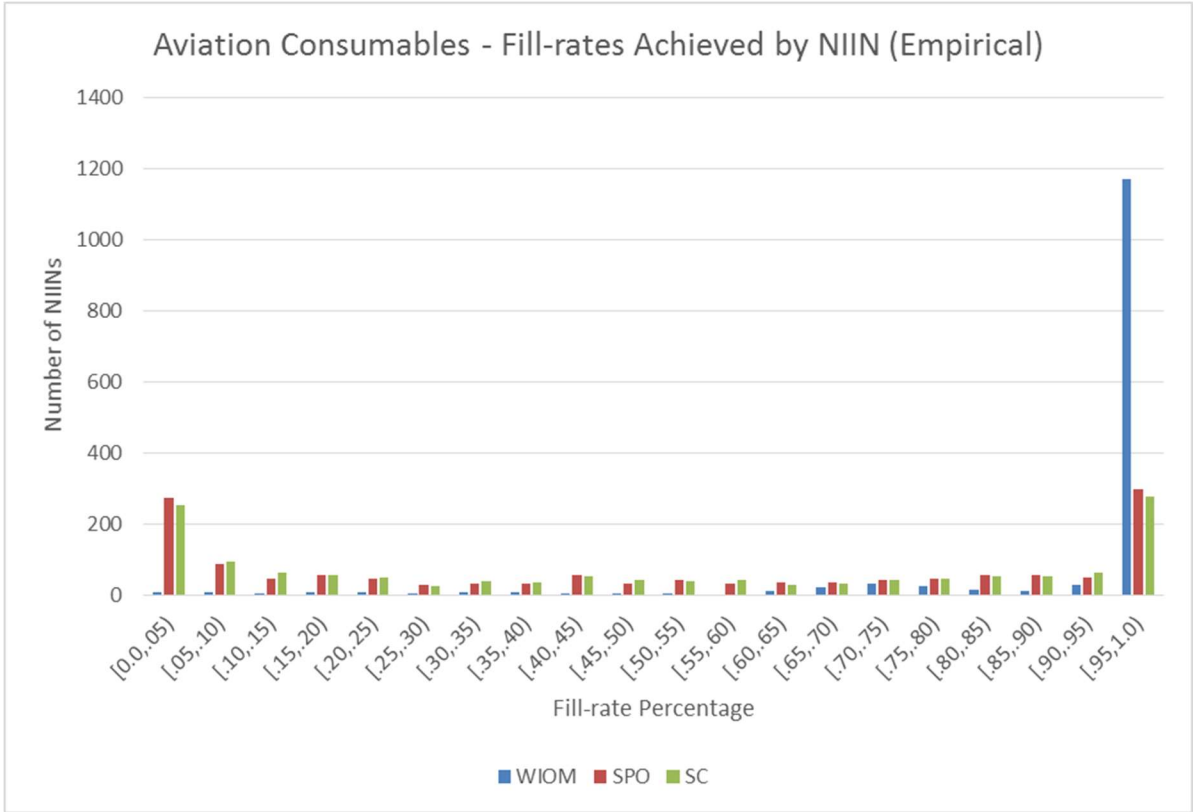
WIOM is more accurate when using the parametric fit than it is using the empirical distributions. The smallest difference between the WIOM estimate and the simulated fill rate was found to be 4.3% in the aviation consumable case but no worse than 16% in any case. The simulated fill rates achieved by WIOM are higher than those achieved by either SPO or SC across all 12 cases. As with the empirical distribution, the fill-rate standard deviation observed between replications is small, indicating that steady state is achieved for all cases.

2. Fill-rates Achieved by NIIN

While the mean fill-rates give an overall picture of the effectiveness of the different methods, examining the distribution of fill-rates inside of the means gives an additional indication of the performance of the reorder point calculations. For example, do a few high fill-rate NIINs offset several low fill-rate NIINs or are the fill-rate values consistent around the mean? Two aviation data sets are presented to answer this question. The fill-rates achieved by the aviation consumable data set are found in Figure 13, the best performing case. The fill-rates achieved by the aviation repairable data set are

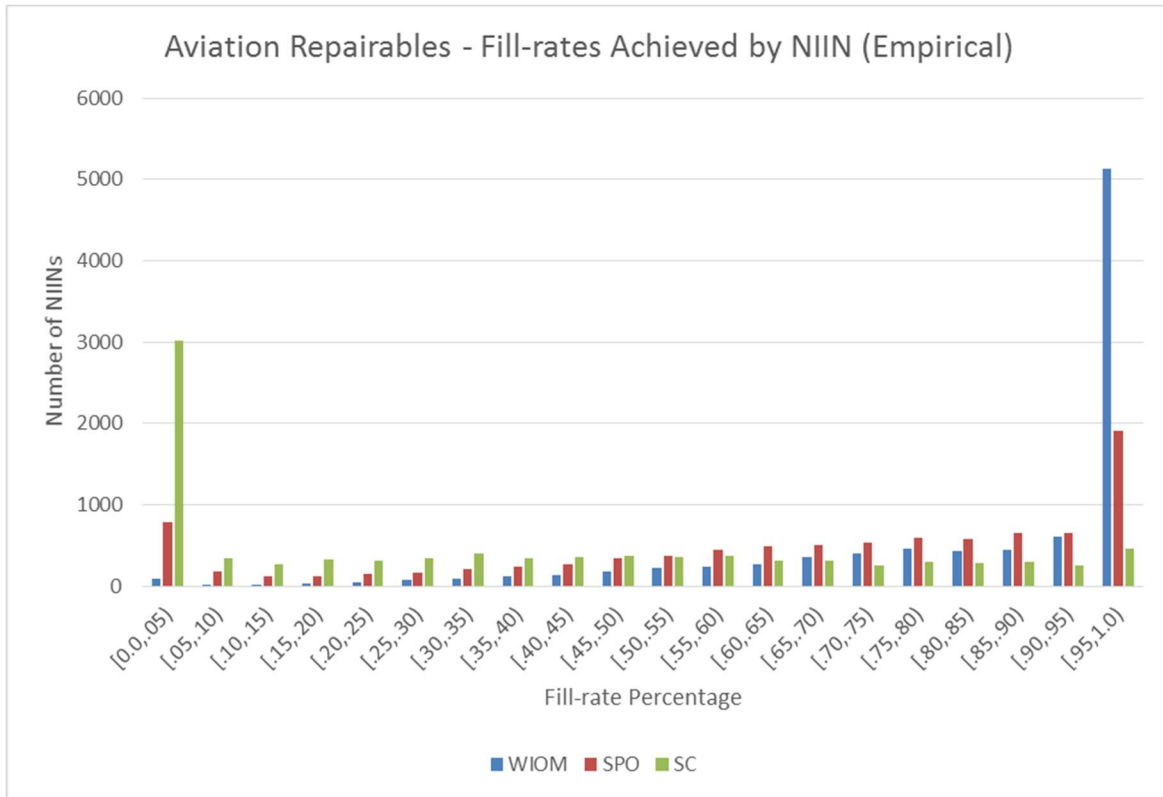
found in Figure 14, the worst performing case. The remaining graphs can be found in Appendix B.

Figure 13. Aviation Consumables - Fill-rates Achieved by NIIN (Empirical)



These are the results by NIIN from the aviation consumable data set using the empirical distribution. The reorder points selected by WIOM achieved very high fill rates with few NIINs in the lower percentages. In contrast, both the SC and SPO had a number of poor performing NIINs offset by an equal number of high performing NIINs.

Figure 14. Aviation Repairables - Fill-rates Achieved by NIIN (Empirical)



These are the results by NIIN from the aviation repairable data set using the empirical distribution.

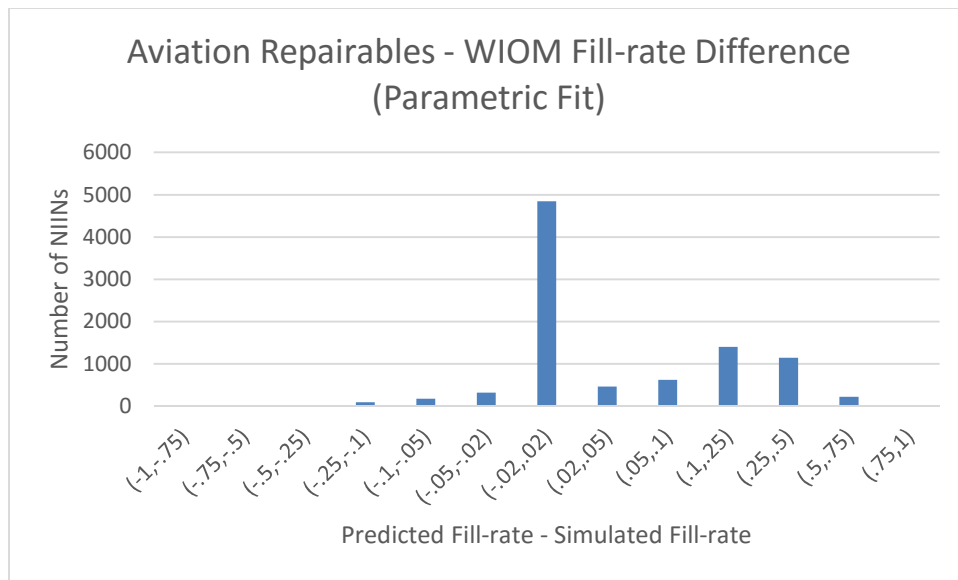
It is important to note that the target fill rate for the NIINs in the data sets are 95%. Any NIINs with a fill rate exceeding 95% is not desirable if it comes at the cost of another NIIN failing to reach 95%. In aviation consumables data set (Figure 13), WIOM is able to achieve the fill-rate goal while minimizing the number of NIINs below the 95% threshold. Both SPO and SC overbuy safety stock for a few NIINs causing them to exceed 95% but then have many more NIINs that fail to reach that threshold.

As shown in Figure 14, WIOM experiences a similar phenomenon in the aviation repairable data set. One possible explanation for the difference in behavior between the two data sets may be the price of the NIINs. The average cost of a NIIN from the aviation consumable data set is \$3,066 while the average cost of a NIIN from the aviation repairable data set is \$80,144. It may not be possible for any of the methods to fit the more expensive items into the safety stock budget.

3. Selection of Demand Probability Distributions

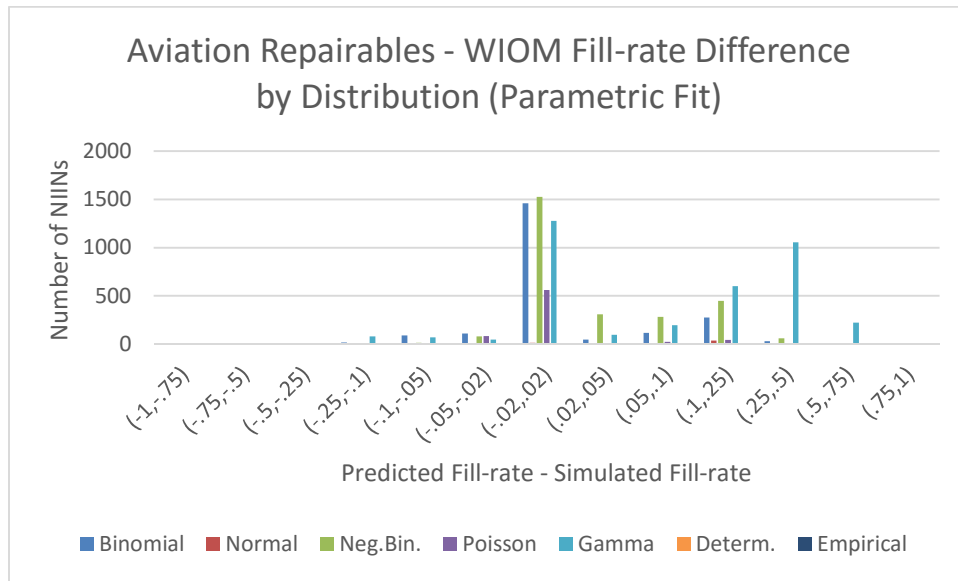
When using the parametric distribution, a probability distribution function is assigned to a particular NIIN based on a parametric fit with simple rules that take the NIIN's demand mean and standard deviation as input. Since there are many different distributions that could be assigned that have the same mean and standard deviation, it is import to check to see if any of the distributions are causing WIOM to consistently over or under-estimate the fill-rate for a NIIN. The graphs of the distribution of fill rates for the aviation repairable data set are found in Figures 15 and 16. The graphs for the remaining data sets are located in Appendix C.

Figure 15. Aviation Repairables - WIOM Fill-rate Difference for All Distributions



The majority of WIOM's estimated fill rates are within 2% of the simulated fill rates. However, when an estimate and simulated value do not match, it is usually because WIOM over-estimated the fill rate.

Figure 16. Aviation Repairables - WIOM Fill-rate Difference by Distribution



When the Poisson distribution can be used in a parametric fit, the resulting estimated fill rate is very close to the simulated value. The binomial, normal and negative binomial also do well. The gamma distribution is less accurate than the other distributions

The graphs show similar trends across the datasets. They all have a right skew demonstrating WIOM’s tendency to over-estimate the fill-rates. The NIINs with demands modeled by the Poisson distribution are the most accurate distribution with the smallest difference between estimated and simulated fill-rates. The NIINs modeled by the gamma distribution show the highest tendency to be least accurately estimated.


C. AVERAGE TIME BACKORDERED

While fill rate is the NAVSUP preferred measure of effectiveness, there are some limitations to its use as a metric. The most important limitation is that all failures to fill a demand are treated equally. That is to say, the item is either available for issue or it is not. The fill-rate for an item does not measure how long it will take to fill a backordered item. As a result, a demand that could be filled by a reorder arriving tomorrow is penalized equally to a demand that cannot be satisfied for one year. This limitation leads to the introduction another measure of effectiveness: the average time backordered. If an inventory reorder policy leads to a lower average time backordered, the failures to fill an

order at the time of the demand will have less of an impact on the customer than an alternative inventory policy.

To explore the backorder space, the mean backorder time is calculated for each NIIN across the 10 runs. These means are then compared as shown in Figures 17 and 18 and Tables 6 and 7. For both the empirical and the parametric demand distributions, the reorder points selected by WIOM showed significant improvement in the median backorder lengths. This was an unexpected result since WIOM does not make any attempt to optimize backorder time. The following table offers an explanation why this occurred. In each case, WIOM selected more reorder points at a higher quantity than SPO. Regardless of the reason, the higher reorder points caused replenishment orders to be placed earlier in the cycle. Since the lead-time for a replenishment order is constant, the earlier the replenishment order is placed the sooner in the cycle it will arrive. This means any backordered demands will be filled more quickly. Figure 19 provides a pictorial example of this phenomena.

Figure 17. Backorder Length - Empirical



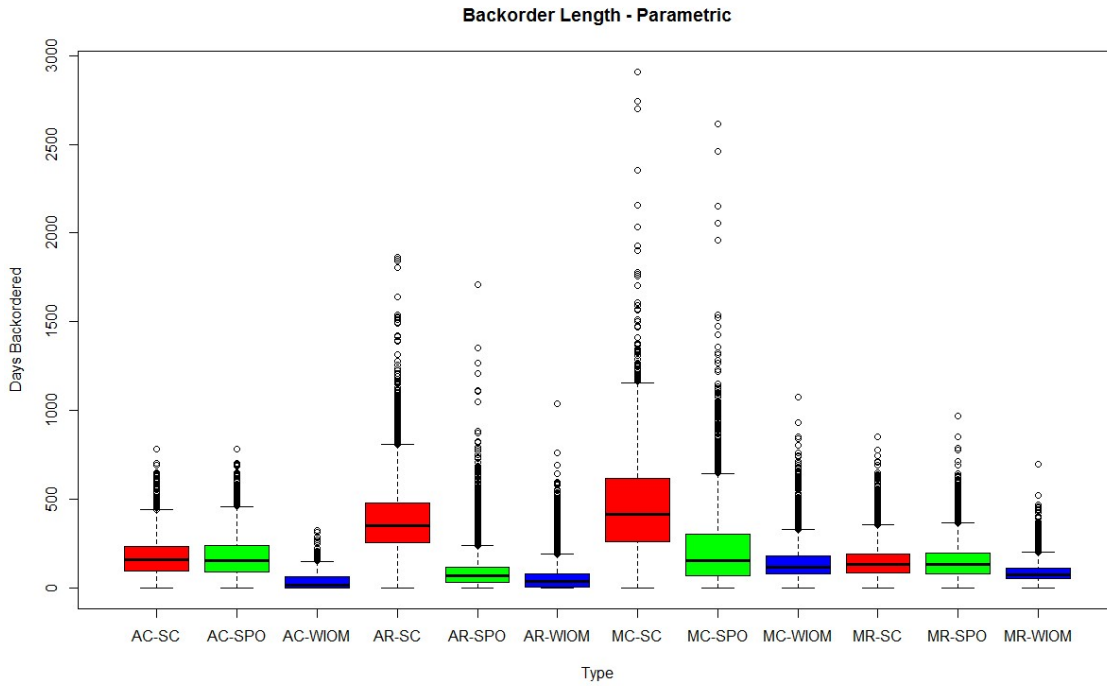
A plot showing the length of backorders in days using reorder points from the three methods using the empirical demand distribution. In all cases, the reorder points selected by WIOM have a lower median backorder length that those using the other two methods.

Table 6. Backorder Length–Empirical

	AC-SC	AC-SPO	AC-WIOM	AR-SC	AR-SPO	AR-WIOM	MC-SC	MC-SPO	MC-WIOM	MR-SC	MR-SPO	MR-WIOM
Minimum	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
First Quartile	93.1	91.4	0.0	254.1	32.3	2.8	260.3	69.6	80.6	83.0	79.6	50.8
Median	158.7	155.6	13.9	351.5	65.4	34.0	414.0	151.9	115.3	132.4	132.9	75.7
Mean	173.8	180.7	36.3	377.9	94.9	57.1	453.9	220.3	150.8	145.9	144.9	87.3
Third Quartile	232.0	238.5	60.8	476.1	115.5	77.3	619.1	300.2	179.5	191.6	194.2	111.5
Maximum	782.4	781.3	323.6	1863.8	1711.0	1038.8	2907.8	2616.3	1072.7	850.6	970.2	694.1

The backorder length in days using reorder points from the three methods using the empirical demand distribution. The reorder points selected by WIOM result in significantly shorter backorder lengths than those selected by SPO and SC.

Figure 18. Backorder Length - Parametric




A plot showing the length of backorders in days using reorder points from the three methods using parametric fit. In all cases, the reorder points selected by WIOM have a lower median backorder length than those using the other two methods.

Table 7. Backorder Length-Parametric

	AC-SC	AC-SPO	AC-WIOM	AR-SC	AR-SPO	AR-WIOM	MC-SC	MC-SPO	MC-WIOM	MR-SC	MR-SPO	MR-WIOM
Minimum	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
First Quartile	88.8	85.8	0.0	28.5	32.2	8.0	91.3	69.9	70.8	82.4	78.9	49.8
Median	154.5	153.6	24.0	61.9	65.1	35.4	178.4	150.7	103.1	131.9	132.7	74.7
Mean	170.9	178.2	38.1	88.4	94.4	58.5	247.7	221.2	140.9	145.4	144.5	86.4
Third Quartile	228.7	233.0	66.8	109.2	114.7	78.4	336.0	305.7	166.1	191.7	193.6	110.7
Maximum	781.5	781.6	244.0	1711.4	1711.4	1052.3	2479.6	2614.2	1054.2	847.7	967.0	712.3

The backorder length in days using reorder points from the three methods using parametric fit. The reorder points selected by WIOM result in significantly shorter backorder lengths than those selected by SPO and the SC.

Figure 19. Backorder Example



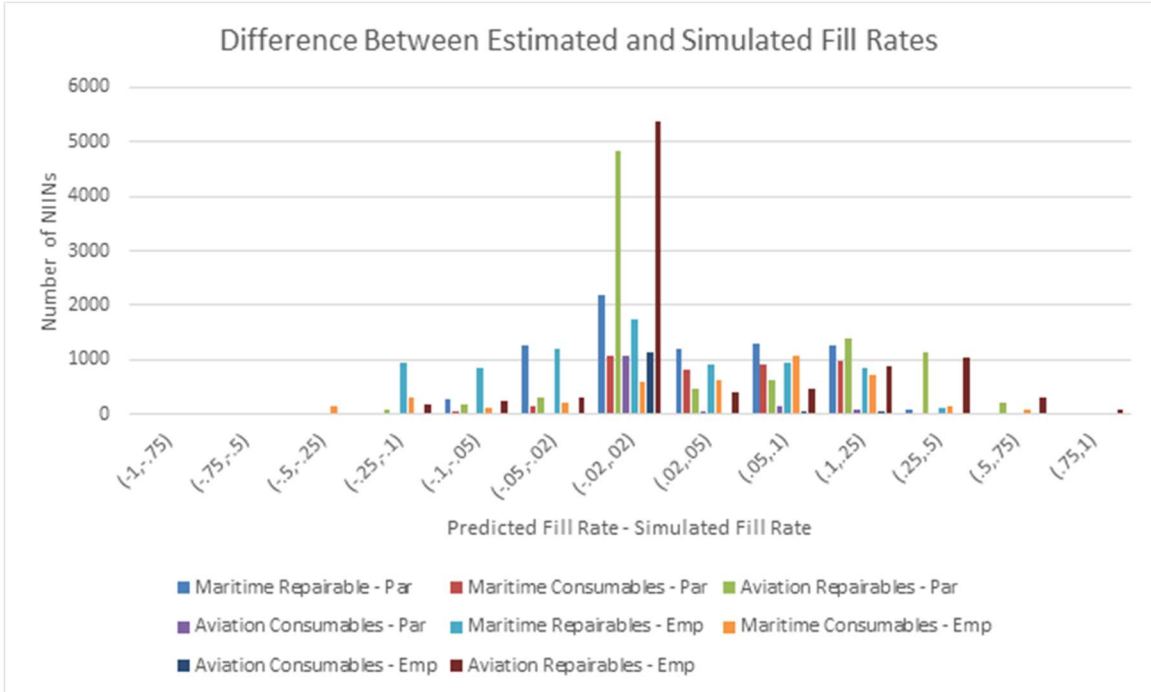
Two different inventory policies are represented. The first has a demand of one unit per time unit, a reorder lead-time of 4, a reorder quantity of 7 and a reorder point of 2. The second has the same characteristics except the reorder point is 1. With the first inventory policy, two units of demand experience backorder; the first for two time units, the second for one time unit. The resulting average time backordered is 1.5 time units. The second inventory policy experiences three backordered units with backorder times of three, two and one time units respectively. The resulting average time backordered is two time units. Since inventory policy one was able to place its order earlier in the cycle, the mean time backordered was lower than that of inventory policy two.

D. DISTRIBUTION ANALYSIS

An important part of WIOM's ability to achieve optimal reorder points is accurately estimating the fill-rate for different NIINs. Figure 20 shows the difference between the fill rate predicted by WIOM and the fill rate simulated by CIS. If the difference between the estimated and simulated fill-rate is positive, the estimate was larger than the simulated value. If the difference between the estimated and simulated fill-rate is negative, the estimate was smaller than the simulated value. The closer the difference is to zero, the more accurately WIOM's estimated fill-rate matched the

simulated fill-rate. As show in Table 8, WIOM was successful in estimating the fill-rates with over 25% of the NIINs falling within 2% of the estimate. Of note, the probability of over-estimating the fill-rate of a NIIN is more than twice the probability of under-estimating the fill-rate.

Figure 20. Difference Between Estimated and Simulated Fill-rates



This figure shows the difference between WIOM’s predicted fill rate and the simulated fill rates for all NIINs across all data sets. The smaller the difference between the two values the closer to the center of the graph the NIIN is plotted.

Table 8. Fill-rate Over-estimation and Under-estimation Across All Datasets

	NIINs	Percentage	Definition
Total	44610		
[-0.02,0.02]	18075	26.40%	Close estimate
(0.02,1.00]	19546	51.67%	Overestimated
[-1.00,-0.02)	6989	21.92%	Underestimated

This table shows the percentage of NIINs across all data sets for which WIOM’s estimated fill rate was a close estimate, and overestimate and an underestimate.

Of particular interest are those NIINs that had estimated values very close to the simulated value and those NIINs that had estimated values very far from the simulated value. Classification trees can be used to identify the characteristics of the NIINs, which could then be used to predict if the NIIN was going to have an estimated fill-rate close to the simulated fill-rate. The aviation repairables dataset was selected for this analysis due to the wide range of differences in estimated and simulated fill-rates. For the remainder of this section we will refer to this difference as fill-rate error.

Two different classification trees are explored. Abbreviations used in the classification trees are listed in Table 9. The first classification tree (Figure 21) uses five groups to categorize the magnitude of the fill-rate error. The results of this tree are listed in Table 11. The second classification tree (Figure 22) divides fill-rate error into two groups (Table 12) to separate low error rates from high error rates. The results of this classification are listed in Table 13.

The five factor classification tree is successful in categorizing 75.61% of the NIINs. It is interesting to note that as the variance to mean and standard deviation to mean ratios increase, the likelihood of misclassification increases. Additionally, the distribution used to describe the demand is only a factor in the final split between groups 1 and 2.

Similarly, the two factor classification is largely dependent on variance to mean and standard deviation to mean ratios. One difference with the five factor classification tree is that reorder time is selected as a factor where a shorter lead times were classified

as having more accurate fill rate predictions. The successful classification rate using this tree is 89.15%.

Table 9. Abbreviation Used in Classification Trees

Abbreviation	Meaning
DemandMe	Monthly mean demand
Distribu	Probability distribution function for demand
DSDC.DMC	Cyclic demand standard deviation / cyclic mean demand
DSDM.DMM	Monthly demand standard deviation / monthly mean demand
DVC.DMC	Cyclic demand variance / cyclic mean demand
Mean	Lead-time mean
ReorderT	Number of days for a reorder to arrive
Var.Mean	Lead-time variance / lead-time mean
Variance	Lead-time variance

These are the abbreviations used in the classification trees in Figure 21 and Figure 22.

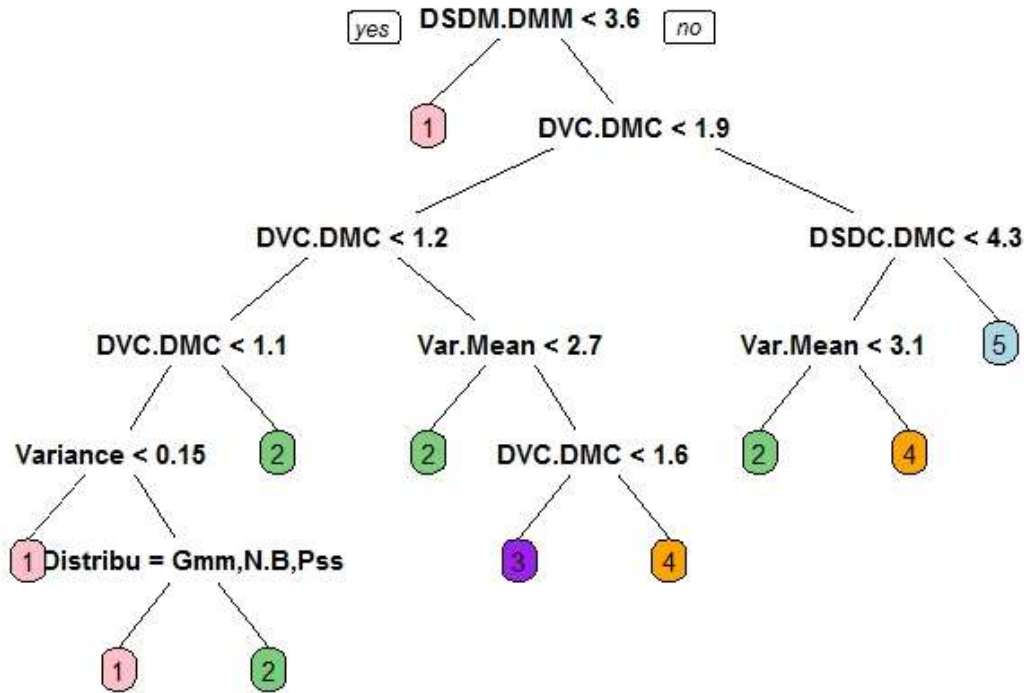
Table 10. Classification Groupings

	Group	Range
Five Factor	1	[0,0.02]
	2	(0.02,0.10]
	3	(0.10,0.20]
	4	(0.20,0.30]
	5	(0.30,1]
Two Factor	1	[0,0.10)
	2	[0.10,1]

This table defines the groups used in the two- and five-factor classification trees.

Figure 21. Five-Factor Classification Tree for Aviation Repairables

Aviation Repairables - Empirical



This classification tree uses characteristics of the NIIN to predict the difference between WIOM's estimated fill rate and the simulated fill rate.

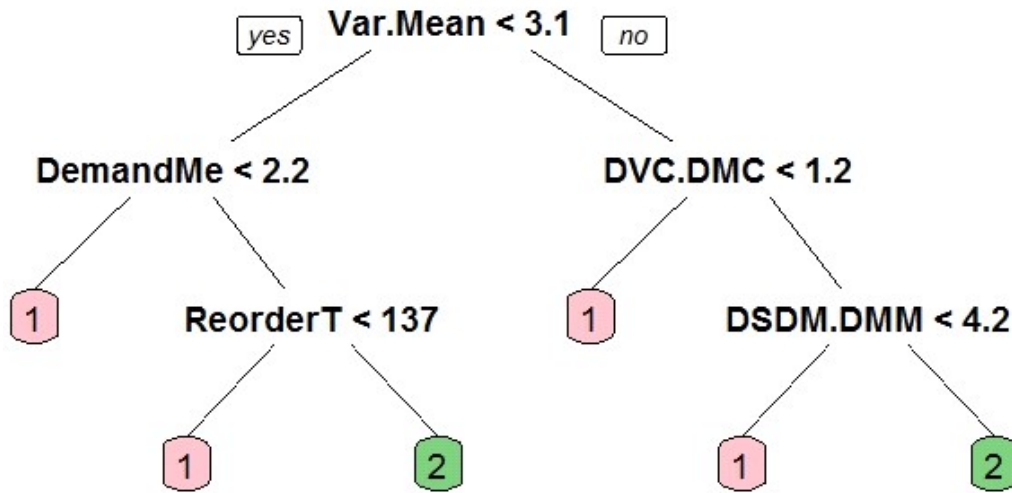
Table 11. Five Factor Classification Tree Results

		Predicted				
		Group	1	2	3	4
Actual	1	5064	286	1	6	0
	2	665	667	39	40	6
	3	418	120	203	77	19
	4	149	10	27	223	73
	5	292	0	0	42	879

If the actual fill-rate error group that a NIIN is categorized in is the same as the predicted fill-rate error group (e.g., on the diagonal) a correct classification has been made. Otherwise, a misclassification has occurred. The classification tree has a misclassification rate of 24.39%.

Figure 22. Two Factor Classification Tree for Aviation Repairables

Aviation Repairables - Empirical



This classification tree uses characteristics of the NIIN to predict the difference between WIOM's estimated fill rate and the simulated fill rate.

Table 12. Two Factor Classification Tree Results

		Predicted	
		Group 1	Group 2
Actual	1	6488	286
	2	724	1808

If the actual fill-rate error group that a NIIN is categorized in is the same as the predicted fill-rate error group (e.g., on the diagonal) a correct classification has been made. Otherwise, a misclassification has occurred. The classification tree has a misclassification rate of 10.85%.

E. SENSITIVITY TO -TIME

The mean amount of time to receive a reorder was the only lead-time information provided by NAVSUP. Since the lead-time is the determining factor as to when a reorder arrives, having some variability in the lead-time may affect the observed fill-rates. To test the potential sensitivity to the factor, the empirical aviation consumable dataset was selected. The lead-time was assumed to be normally distributed with a mean as provided by NAVSUP and different levels of variance. Table 11 shows the results of this analysis. While there was a slight decrease in the simulated fill-rate as variability increased, the overall effect was negligible. However, assuming that the normal distribution applies in this case may not be correct. Additional information about the true nature of the lead-time distribution may lead to different results.

Table 13. Sensitivity of the Fill-rate to Stochastic Lead-time

Distribution	Std Dev as % of Mean	Simulated Fill Rate
Normal	0	0.9432
Normal	10	0.9431
Normal	20	0.9430
Normal	30	0.9430

This chart shows the sensitivity of the fill rate to various stochastic lead-times. Under the normal distribution, there is little change in the simulated fill rate. Other distributions may yield different results but no data was available to suggest other distributions or their parameters.

THIS PAGE INTENTIONALLY LEFT BLANK

V. CONCLUSIONS AND RECOMMENDATIONS FOR FOLLOW-ON STUDIES

This thesis provides qualitative and quantitative information to help NAVSUP explore the effects of using WIOM, SPO and the SC to set reorder points across a variety of NIINs. This section summarizes the findings and provides some recommendations for future work.

A. CONCLUSIONS

The overarching goal of CIS is to produce a model that produces an estimated fill-rate based on a demand probability distribution and various parameters that describe the NIIN. CIS is able to draw a number of different conclusions about the three different methods for determine a reorder point:

- When considered as an entire system, WIOM's fill-rate estimations range between 1 and 18% of the simulated values. WIOM's fill-rate estimations are closer to the simulated value when the parametric fit is used.
- The reorder points selected by WIOM resulted in higher average fill-rates than those selected by SPO and the SC.
- SPO's tendency to select higher reorder points for high demand items had two consequences. The first is a higher overall fill-rate when compared with WIOM in some cases. The second is a higher mean backorder length when compared with WIOM.
- The SC consistently underperformed both WIOM and SPO.
- A demand probability distribution fit to a gamma probability distribution function is more likely to have a large deviance between the WIOM estimated fill-rate and the simulated fill-rate than any other probability distribution. The gamma function should be avoided whenever possible.

One of the advantages of CIS is that it offers a level playing field to evaluate the different methods for finding the reorder point. The only difference between simulations of the different methods was the reorder points themselves; how those reorder points were determined was unimportant and ignored by CIS. This provides an unbiased comparison of the three methods and can help to inform a decision regarding future reorder point setting algorithms.

Another advantage of CIS is the ability to calculate backorder statistics for comparison between the three methods. This ability was not previously available in previous analysis and offered additional insight into the differences in the outcomes between SPO and WIOM.

B. RECOMMENDATIONS FOR FOLLOW-ON RESEARCH

While this thesis sought to explore many different aspects of the reorder point establishment problem, there are many opportunities for follow-on research. This thesis briefly explored the sensitivity of the WIOM selected reorder points to variability in lead-time, a more representative distribution could be found which may lead to different results. Additionally, sensitivity to other input parameters could be explored.

One of the assumptions of the simulation is that demands for items are always for a single quantity. While this is true for the repairable datasets, it is usually a false assumption for the consumable datasets. Relaxing this assumption and using probability distributions for the number of items in a single demand event would allow for a more realistic simulation.

This thesis considered NIINs collectively regardless of their individual properties. NIINs with very low demand and high variance were grouped with NIINs with very high demand and low variance. WIOM has an option to use level setting strategy indicators (LSSI) where NIINs of similar properties are grouped together. Future work could consider the characteristics of a NIIN that should put it in one LSSI or another. Additionally, a comparison could be made between the results of fill rates achieved using LSSIs with the fill rates of the single category of NIINs.

Lastly, both WIOM and the simulation are heavily dependent upon the probability distributions for the demand. The parametric fit for the demand distribution function relied solely on the lead-time mean and standard deviation. Future work could examine these fits to determine if they could be refined with additional information about the NIIN.

APPENDIX A. SUPPLY TERMS

Backorder—When a demand for a NIIN cannot immediately be satisfied due to insufficient stock on hand, the request for the item is placed into a queue for fulfillment upon receipt of a replenishment order. A demand in this state is said to be backordered.

Cycle—A measure of the amount of time defined by $t_i' = \frac{t_i}{\max\left(1, \frac{\hat{x}_i}{Q_i}\right)}$ where t_i is

the lead-time, \hat{x}_i is lead-time demand and Q_i is the reorder size.

Demands—A request from an end user for an item. In this thesis, all demands are assumed to be for a single unit of the item.

Fill-rate—The percentage of demands that can be immediately be satisfied by an issue from on-hand inventory.

Inventory position—The sum of the number of items in inventory plus the number of items on order minus the number of items backordered.

Lead-time—The amount of time between an order being placed for additional stock and the receipt of those items.

NIIN—National identification number. A 9-digit subset of the NSN.

Reorder point (s)—The answer to the question of when to order additional units for inventory. When the inventory position reaches the reorder point, an order is placed for additional reorder quantity of the item, which will be delivered at the end of the lead-time.

Reorder quantity (Q)— The amount of an item ordered during a reorder.

Safety stock—Extra inventory kept on hand to act as a hedge in case actual demand exceeds expected demand during the lead-time. A backorder occurs when a demand is processed and safety stock is exhausted.

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX B. FILL-RATE GRAPHS BY DATASET

Figure 23. Maritime Consumables: Fill-rates Achieved by NIIN (Empirical)

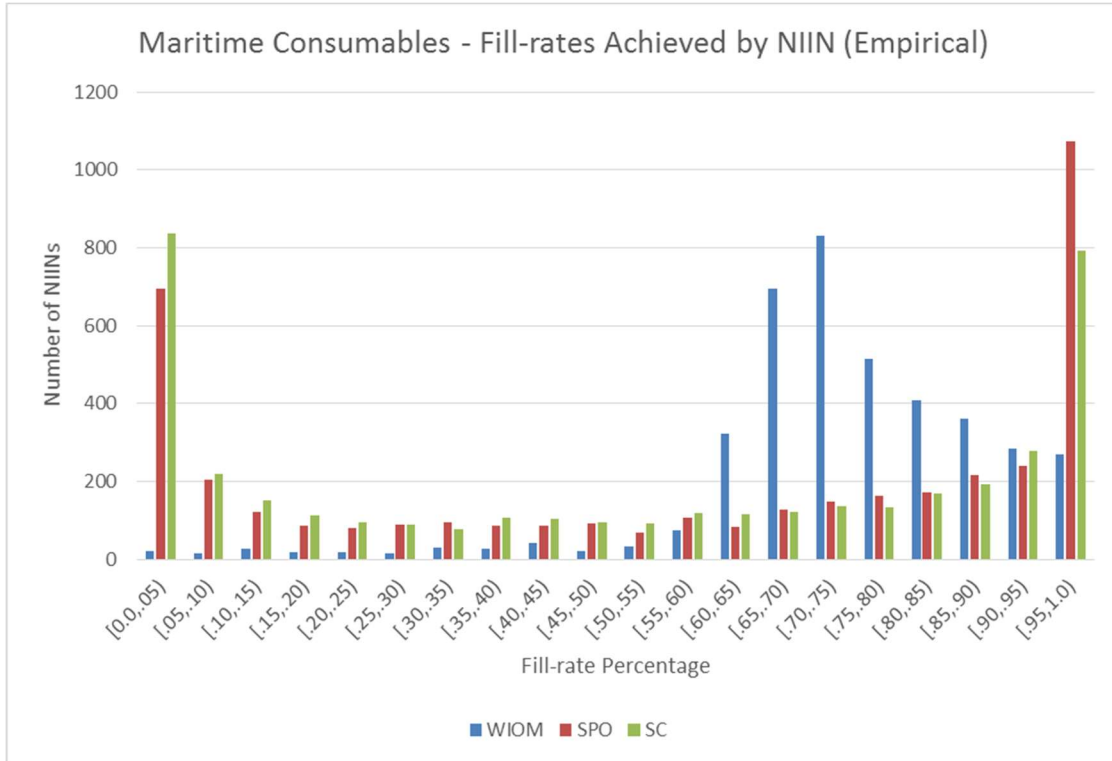


Figure 24. Maritime Repairables: Fill-rates Achieved by NIIN (Empirical)

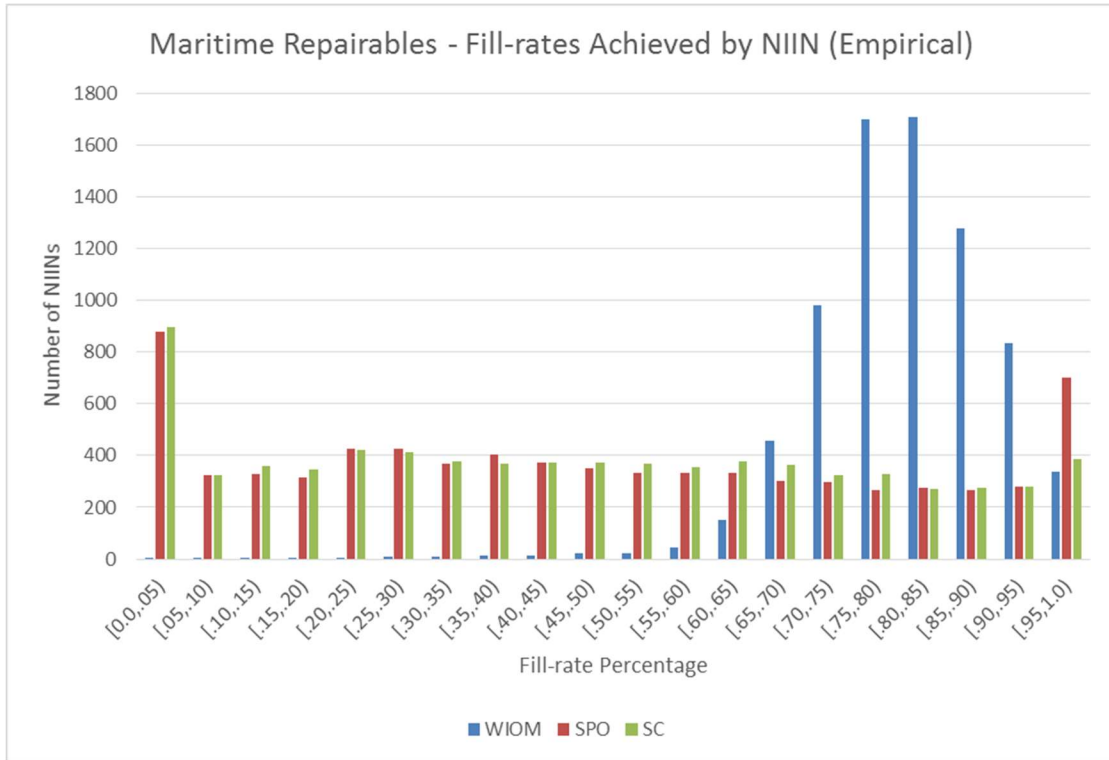


Figure 25. Aviation Consumables: Fill-rates Achieved by NIIN (Parametric)

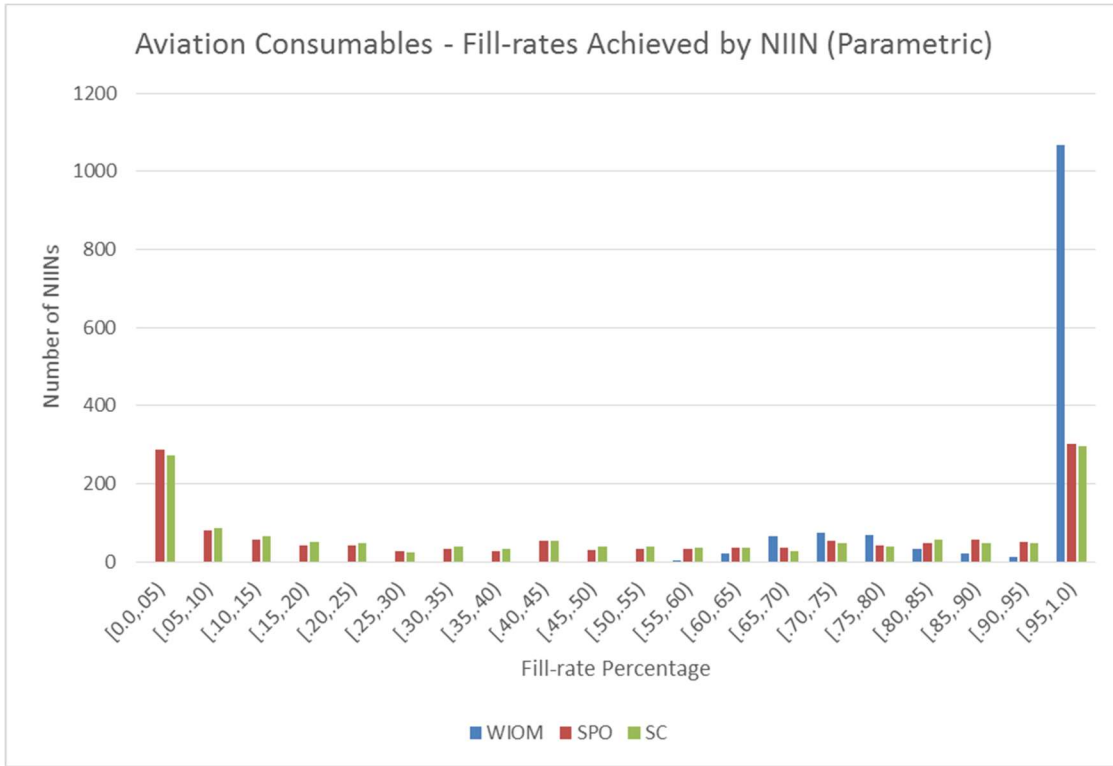


Figure 26. Aviation Repairables: Fill-rates Achieved by NIIN (Parametric)

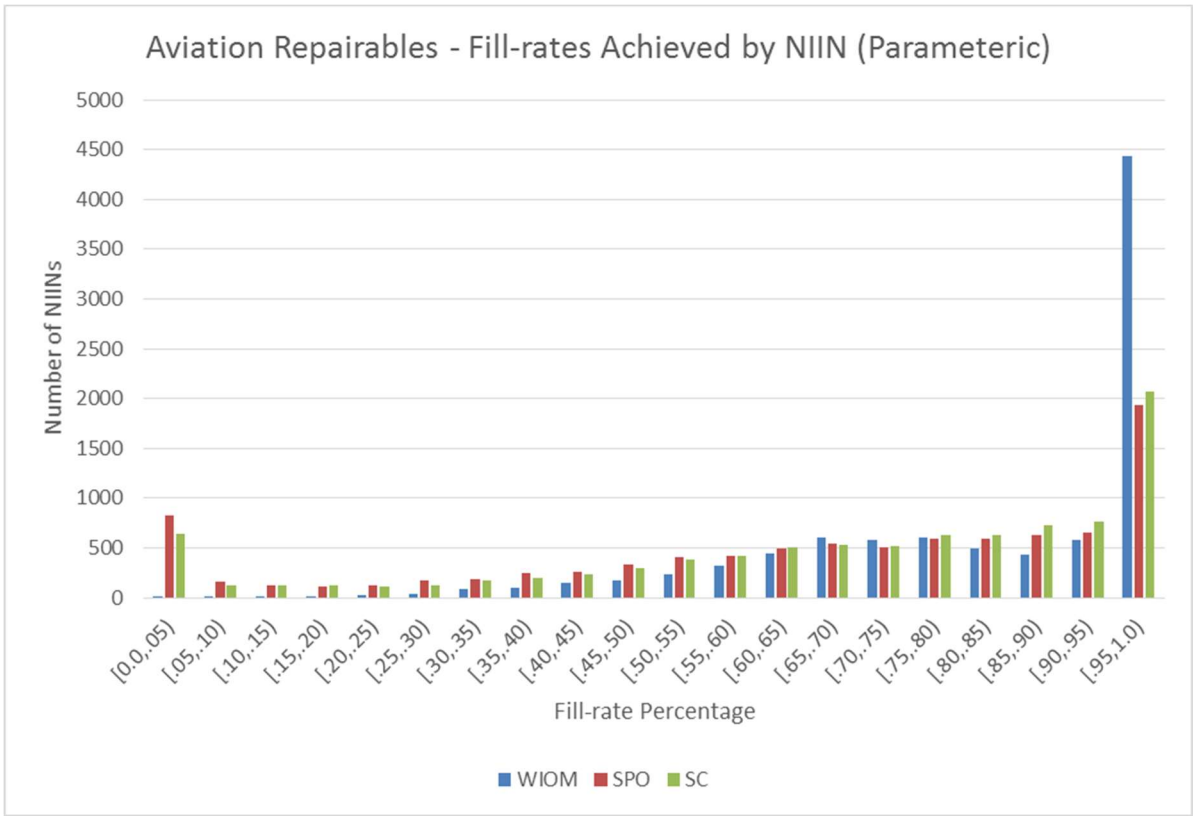


Figure 27. Maritime Consumables: Fill-rates Achieved by NIIN (Parametric)

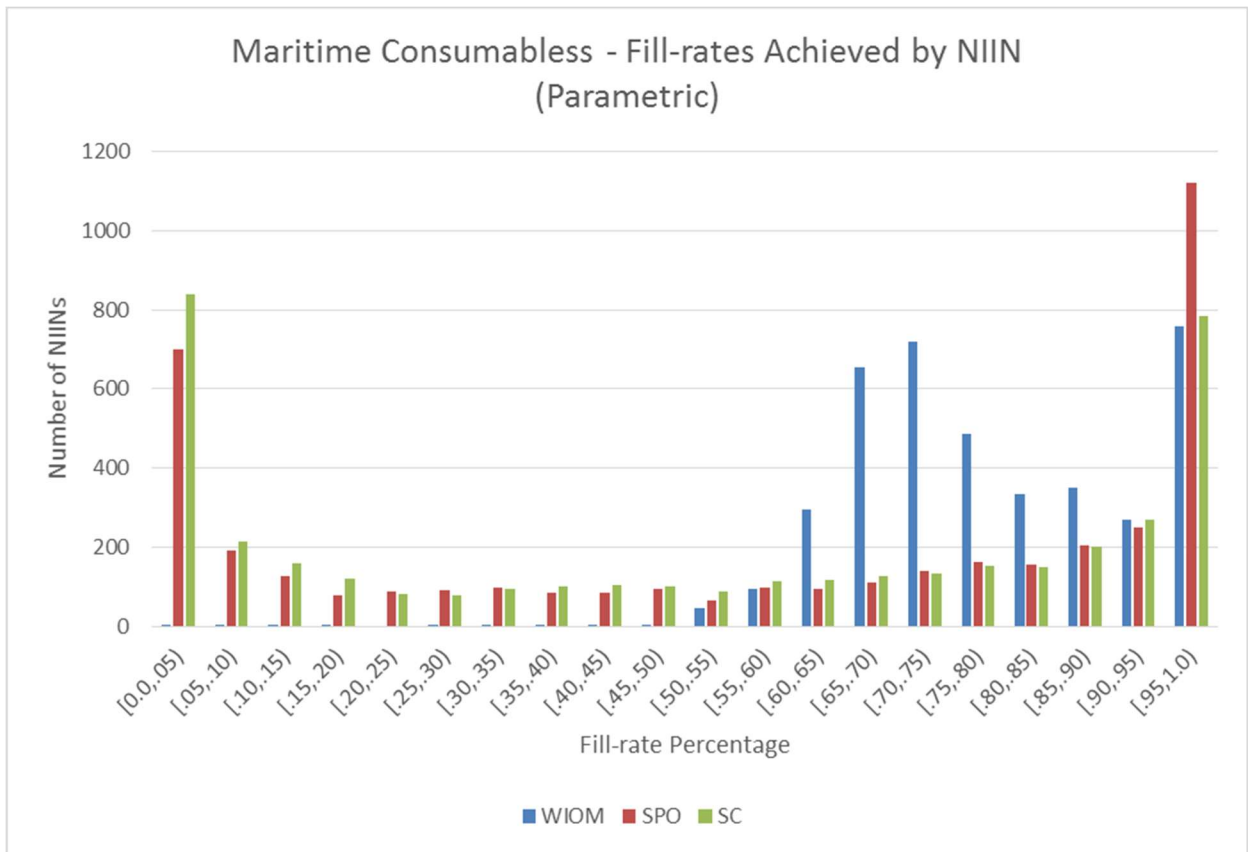
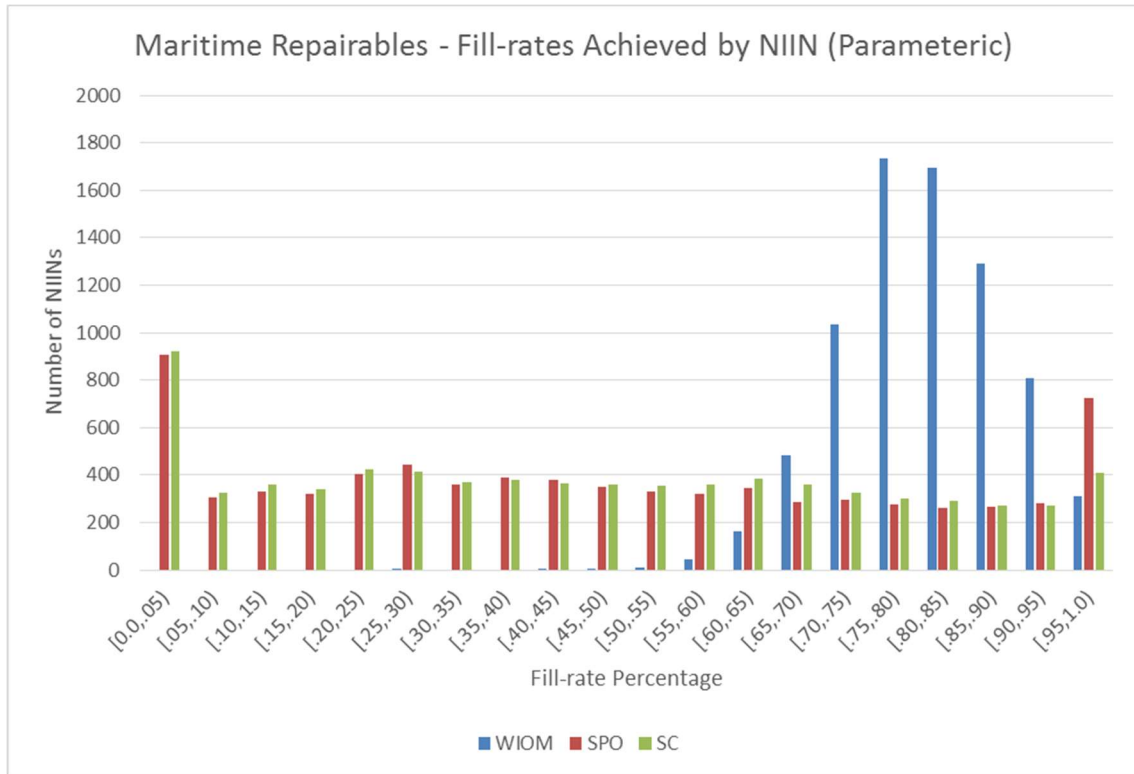


Figure 28. Maritime Repairables: Fill-rates Achieved by NIIN (Parametric)



APPENDIX C. FILL-RATE DISTRIBUTION GRAPHS

Figure 29. Aviation Repairables: WIOM Fill-rate Difference for All Distributions

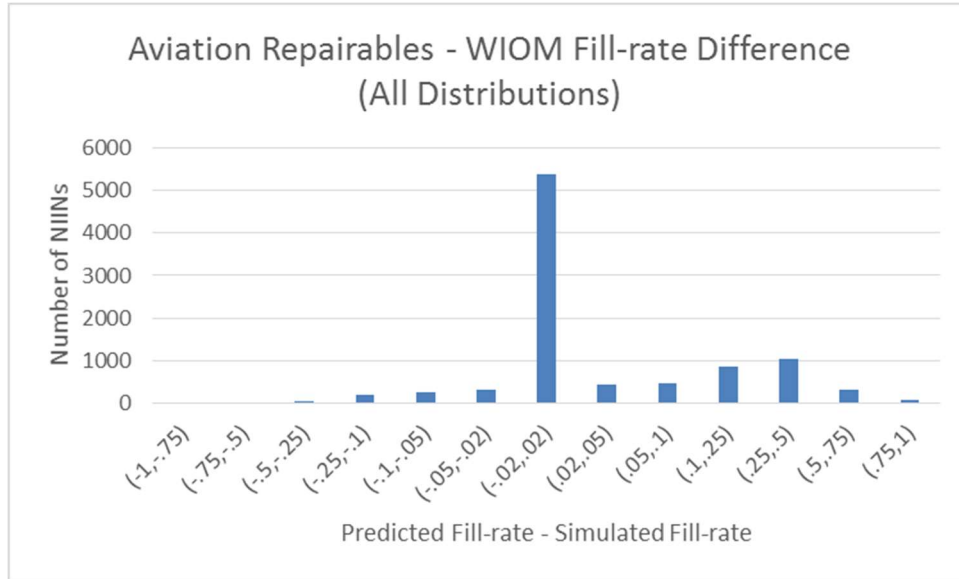


Figure 30. Aviation Repairables: WIOM Fill-rate Difference by Distribution




Figure 31. Maritime Consumables: WIOM Fill-rate for All Distributions

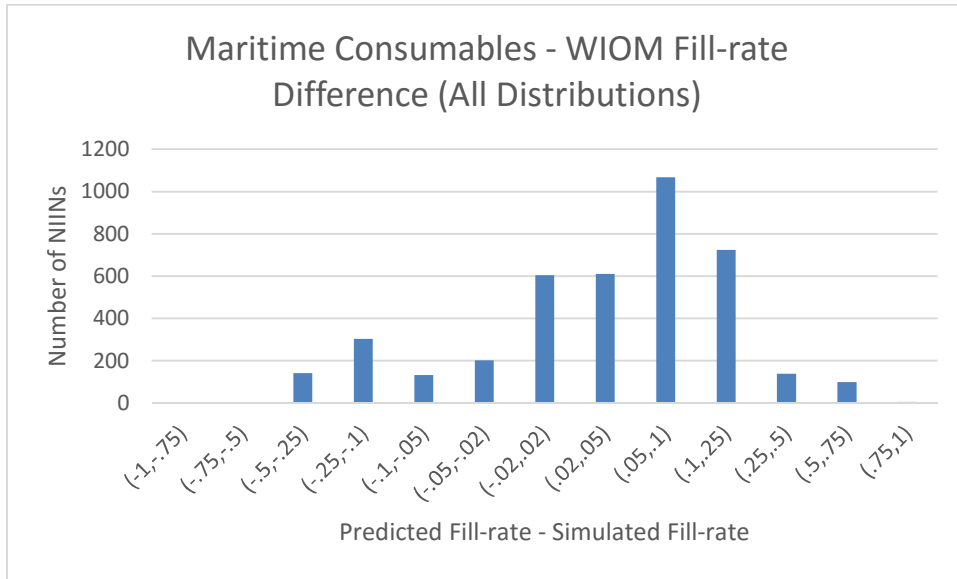


Figure 32. Maritime Consumables: WIOM Fill-rate by Distribution

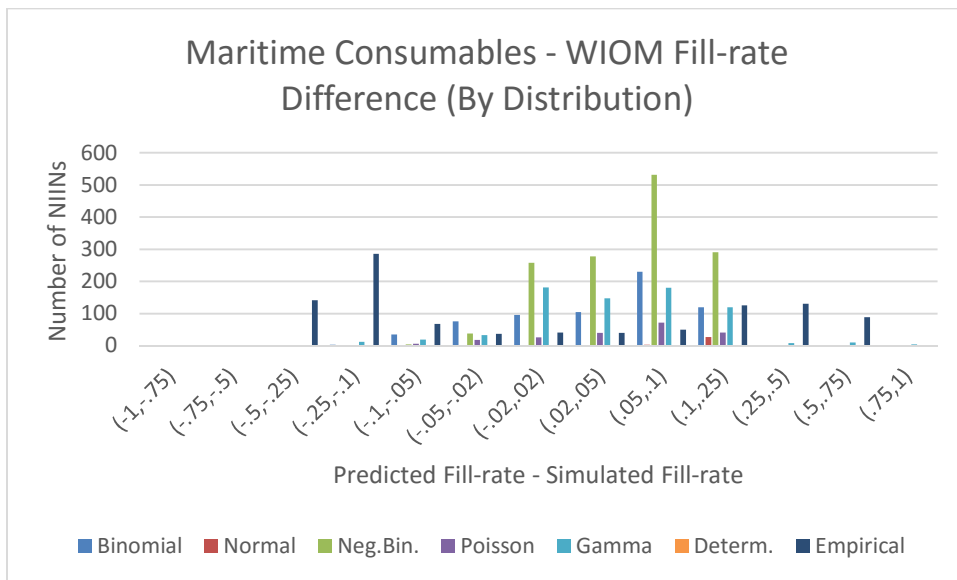


Figure 33. Maritime Repairables: WIOM Fill-rate Difference for All Distributions


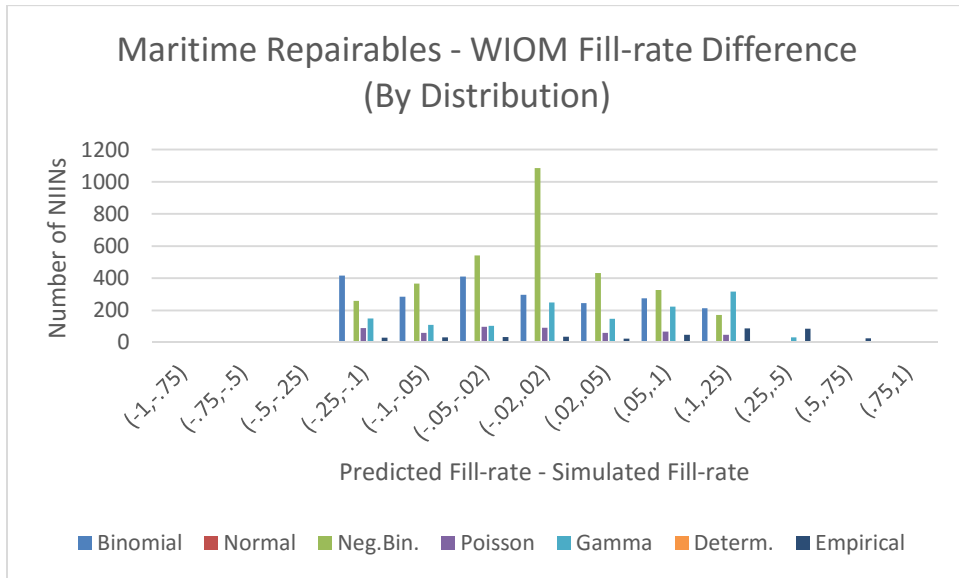


Figure 34. Maritime Repairables: WIOM Fill-rate Difference by Distribution



THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF REFERENCES

- Buss AH (1996) Modeling with Event Graphs. Charnes JM, Morrice DJ, Brunner DT, Swain JJ, eds. *Proceedings of the 2006 Winter Simulation Conference*. (Institute of Electrical and Electronics Engineers, Washington, DC), 153-160.
- Cioppa TM, Lucas TW (2007) Efficient nearly orthogonal and space-filling Latin Hypercubes. *Technometrics* 4 (1): 45–55.
- Conway RW (1963) Some Tactical Problems in Digital Simulation. *Management Science* 1 (1): 47–61.
- Law AM, Kelton WD (2007) *Simulation Modeling and Analysis*, 4th ed. (McGraw-Hill, New York)
- Office of the Chief of Naval Operations (2011) OPNAVINST 4442.5A: Readiness Based Sparing. Instruction, Office of the Chief of Naval Operations, Arlington, VA.
- Oswald, AJ (2014) Electronic message between author and LCDR Andrew Oswald, October 21.
- Rardin, RL (1998) *Optimization in Operations Research* (Prentice Hall, Upper Saddle River, New Jersey).
- Salmeron J, Craparo E (2014) Wholesale Inventory Optimization Models. Working paper, Naval Postgraduate School.
- Sanchez PJ (2007) Fundamentals of Simulation Modeling. Henderson SG, Biller B, Hsieh M, Shortle J, Tew JD, and Barton RR, eds. *Proceedings of the 2007 Winter Simulation Conference*. (Institute of Electrical and Electronics Engineers, Washington, DC).
- Sanchez SM (2008) Better Than a Petaflop: The Power of Efficient Experimental Design. Hill R, Monch L, Rose O, Jefferson T, Fowler JW. *Proceedings of the 40th Conference on Winter Simulation*. (Winter Simulation Conference: IEEE, Washington, DC).
- Schruben L (1983). Simulation Modeling with Event Graphs. *Communications of the ACM* 26 (11): 957 – 963.
- Silver EA, Pkye DF, Peterson R (1998) *Inventory Management and Production Planning and Scheduling*. (John Wiley & Sons, New York).
- Welch, PD (1983). The Statistical Analysis of Simulation Results. In *The Computer Performance Modeling Handbook* (Academic Press, New York).

THIS PAGE INTENTIONALLY LEFT BLANK

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
Ft. Belvoir, Virginia
2. Dudley Knox Library
Naval Postgraduate School
Monterey, California