NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

SENSITIVITY ANALYSIS OF DEMAND DISTRIBUTION FOR NAVAL AVIATION READINESS-BASED SPARING MODEL

by

Paul M. Cabe

September 2020

Thesis Advisor:	Arnold H. Buss
Co-Advisor:	Javier Salmeron-Medrano
Second Reader:	Duncan Ellis (U.S. Navy)

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** September 2020	**3. REPORT TYPE AND DATES COVERED** Master's thesis		
4. TITLE AND SUBTITLE SENSITIVITY ANALYSIS OF DEMAND DISTRIBUTION FOR NAVAL AVIATION READINESS-BASED SPARING MODEL				**5. FUNDING NUMBERS**
6. AUTHOR(S) Paul M. Cabe				
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.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release. Distribution is unlimited.				**12b. DISTRIBUTION CODE** A

13. ABSTRACT (maximum 200 words)

Readiness Based Sparing (RBS) models support the life cycle of any system through the optimization of stock allowance levels. Optimal RBS results are essential to maintain fleet readiness at an acceptable cost. Naval Aviation RBS Model (NAVARM) is the tool used by Naval Supply Systems Command to plan the stock allowances for embarked airwings and shore-based aircraft. In order to gain confidence in NAVARM results, it is necessary to validate some modeling assumptions that have not been tested to date. RBS models like NAVARM assume that the distribution of the mean time between failures (MTBF) for any part is exponential. This assumption may not hold in practice for certain parts. Therefore, a question arises as to whether the quality (operational availability by cost) of the solution provided by NAVARM is subject to the effects of this assumption.

This thesis tests the influence of the MTBF distribution on operational availability using the Readiness-Based Sparing Simulation (RBSIM) developed by a former Naval Postgraduate School student. We test the alternate distributions Weibull, gamma, and log-normal, with mean to variance ratios (MTVRs) of 1.5 and 0.5. These MTBF distributions are applied to either all parts or a select subset of parts (based on demand). Initial results on the aviation consolidated allowance list for the USS Carl Vinson (CVN 70) show that both distribution type and MTVR may have a significant effect on operational availability of all weapon systems.

14. SUBJECT TERMS Readiness Based Sparing model, sensitivity analysis, NAVARM			**15. NUMBER OF PAGES** 75
			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

NSN 7540-01-280-5500 Standard Form 298 (Rev. 2-89) Prescribed by ANSI Std. 239-18

THIS PAGE INTENTIONALLY LEFT BLANK

SENSITIVITY ANALYSIS OF DEMAND DISTRIBUTION FOR NAVAL AVIATION READINESS-BASED SPARING MODEL

Paul M. Cabe
Lieutenant Commander, United States Navy
BSBA, Auburn University, 2009

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN OPERATIONS RESEARCH

from the

NAVAL POSTGRADUATE SCHOOL
September 2020

Approved by: Arnold H. Buss
Advisor

Javier Salmeron-Medrano
Co-Advisor

Duncan Ellis
Second Reader

W. Matthew Carlyle
Chair, Department of Operations Research

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

Readiness Based Sparing (RBS) models support the life cycle of any system through the optimization of stock allowance levels. Optimal RBS results are essential to maintain fleet readiness at an acceptable cost. Naval Aviation RBS Model (NAVARM) is the tool used by Naval Supply Systems Command to plan the stock allowances for embarked airwings and shore-based aircraft. In order to gain confidence in NAVARM results, it is necessary to validate some modeling assumptions that have not been tested to date. RBS models like NAVARM assume that the distribution of the mean time between failures (MTBF) for any part is exponential. This assumption may not hold in practice for certain parts. Therefore, a question arises as to whether the quality (operational availability by cost) of the solution provided by NAVARM is subject to the effects of this assumption.

This thesis tests the influence of the MTBF distribution on operational availability using the Readiness-Based Sparing Simulation (RBSIM) developed by a former Naval Postgraduate School student. We test the alternate distributions Weibull, gamma, and log-normal, with mean to variance ratios (MTVRs) of 1.5 and 0.5. These MTBF distributions are applied to either all parts or a select subset of parts (based on demand). Initial results on the aviation consolidated allowance list for the USS Carl Vinson (CVN 70) show that both distribution type and MTVR may have a significant effect on operational availability of all weapon systems.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF FIGURES

LIST OF TABLES

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS AND ABBREVIATIONS

ARROWS	Aviation Readiness Requirements Oriented to Weapons Replaceable Assemblies
A_o	operational availability(ies)
AVCAL	aviation consolidated allowance list
CNO	Chief of Naval Operations
EBO	expected backorder
FIFO	first in first out
LRU	line replaceable unit
MALS	Marine Aviation Logistics Squadrons
MTBF	mean time between failure
MTTR	mean time to repair
MTVR	mean to variance ratio
NAVARM	Navy Aviation Readiness Based Sparing Model
NAVSUP	Navy Supply Systems Command
NAVSUP, WSS	Naval Supply Systems Command, Weapon Systems Support
NMCS	non-mission capable supply
OPNAVINST	Operational Navy Instruction
PUK	pack-up kit
RBS	Readiness Based Sparing
RBSIM	Readiness-Based Sparing Simulation
SHORECAL	Shore consolidated allowance list
SRA	shop replaceable assembly
SRU	shop replaceable units
WS	weapon system

THIS PAGE INTENTIONALLY LEFT BLANK

EXECUTIVE SUMMARY

Naval Aviation uses a Readiness Based Sparing (RBS) planning model to develop stock levels for aircraft weapon systems (WS) for both afloat platforms and shore facilities. The RBS model known as Naval Aviation RBS Model (NAVARM) is the RBS model Naval Supply Systems Command (NAVSUP) employs to capture the allowances levels needed in order to achieve a directed level of WS operational availability (A_o) targets.

NAVARM is designed as a single-site, multi-indentured model. It uses historical demand from the Navy Supply system along with other internal algorithmic inputs to produce cost-effective allowance levels that are needed to achieve A_o for each WS. NAVARM's allowance levels are provided for all the parts that comprise each WS.

NAVARM assumes an exponential distribution for mean time between failure (MTBF) to build the allowance levels. The purpose of this thesis is to test what would occur if NAVARM's distribution assumption were incorrect; in particular, by hypothesizing the Weibull, gamma, and log-normal distributions as alternates to the default exponential. Like the exponential, these alternate distributions have theoretical support in some cases, as discussed in the literature.

In order to test NAVARM's assumption we employ an existing NAVARM simulation model (RBSIM) developed by Wray at the Naval Postgraduate School. RBSIM uses NAVARM's output allowance as an input, and simulates failures at the individual component level. It accomplishes this task by generating a stochastic failure time of individual parts based on the MTBF distribution for each part. The original RBSIM has been modified in this thesis in order to utilize the alternate distributions as needed. In addition, we also specify a mean to variance ratio (MTVR) of 1.5 and 0.5 for those distributions. RBSIM develops an expected completion time of repair and return to inventory based on the assigned failure. The failure rates are also specific to the part's position within the WS. The simulation only calculates the metrics that relate to readiness in order reduce the run time and complexity.

The scenario tested in this thesis is from the aviation consolidated allowance list (AVCAL) from the USS Carl Vinson (CVN 70). The AVCAL is built using a 24-month demand history of aviation parts. The AVCAL consists of 7 different WSs which represent all the different types of aircraft that currently deploy aboard the carrier. The AVCAL contains approximately 42,000 unique parts. RBSIM simulates failure of the parts first using only alternate distributions; then, with a mixture of the alternates and exponential. The mixed distribution method's intent is to capture results where the possibility that just some parts violate the RBS distributional assumption. The mixed distribution applies the alternate distribution to approximately 16,000, low-demand parts (less than 5 demands over the 24-month period).

Based on the use of single-distribution or mixed-distributions, we generate 13 simulation results. There is a modest discrepancy between the simulation results for the exponential-only calculated A_o and NAVARM's calculated A_o. RBSIM calculates the exponential-only A_o slightly higher than NAVARM. This discrepancy is more likely due to an error in RBSIM calculations than an error in NAVARM. We hypothesize that because NAVARM has been extensively tested by NAVSUP, but have not confirmed it within this research.

Our overall results show that 33% of the alternate-only distributions with MTVR of 1.5 and 100% of the alternate-only distributions with 0.5 MTVR achieve A_o for all WSs. The gamma-only distribution with a 0.5 MTVR, and both the log-normal-only with MTVR values of 1.5 and 0.5, also achieve A_o across all WSs. The Weibull-only distribution with a 1.5 MTVR achieves the lowest A_o among all of the alternate distributions.

The E2-D and H60-R WSs fail to achieve A_o using the alternate-only distributions most often. Both WSs closely resemble each other's A_o for any distribution. The EA-18G WS achieves A_o regardless of distribution or MTVR. The F/A-18 E and F aircraft models differ on A_o results from the EA-18G, even though all three share the same logistical support pipeline. The H-60 R and S aircraft have the greatest difference between themselves for RBSIM's A_o results even though they too share the same logistical support.

In conclusion, it appears that an alternate distribution may have a significant impact on A_o, depending on the distribution type, MTVR and parts that have that MTBF distribution. Therefore, if the distribution of MTBF does not align with the RBS assumed distribution then there are indications that the achieved A_o will not meet model projections.

Although this thesis concentrates on the demand distribution of NAVARM, future analysis of NAVARM's other assumptions is recommended. An in-depth analysis of RBSIM is also warranted given we observe modest unexpected differences with NAVARM's estimated A_o when using the exponential distribution. Lastly, further statistical analysis on actual demand distribution of parts would be beneficial given we have already identified that alternate distributions affect A_o.

THIS PAGE INTENTIONALLY LEFT BLANK

ACKNOWLEDGMENTS

First and foremost, I would like to thank my wife for her undying patience and encouragement throughout this entire process. I could not have done it without your support!

To Professor Salmeron, thank you for everything you have done during the building of this thesis. Not only did you guide me, you took your time to teach me all the details I would need in order to succeed with this analysis.

To Professor Buss, thank you for modifying RBISM’s code in order for us to perform the simulations. Without your input throughout this whole process, I surely would not have been able to begin to tackle this problem.

To Commander Ellis, thank you for this idea. If it weren’t for you, I would not have this thesis topic. Also thank you for agreeing to be my second reader when you did not have to.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

Material readiness of the U.S. Navy is critical to national security. It ensures that naval vessels and aircraft have the material on hand to carry out their mission to protect the U.S. and its interests. A memo by Secretary of Defense James Mattis in 2018 to the Navy and Marine Corps requires both services to immediately improve the operational availability of the F/A-18 Hornet fighter jet to meet the minimum operational readiness goal of 80%. (J. Mattis, personal correspondence, 17 Sep 2018). The main issue he addresses is the need to improve maintenance practices and have the correct parts readily available. Navy Supply Systems Command (NAVSUP) is the lead organization responsible for developing and implementing the supply chain processes to accomplish this goal.

NAVSUP uses readiness requirements that are established by the fleet to build material support strategies for naval aviation by the most economical and efficient means possible, including the use of planning models. The Navy's planning model for logistical support at the retail level is known as the Readiness Based Sparing (RBS) model. RBS allows the Navy to build parts allowances for organizational and intermediate level maintenance for aircraft that meet required availability at a low, potentially optimal cost.

The Navy Aviation Readiness Based Sparing Model (NAVARM) is a planning model currently in use by Naval Supply Systems Command, Weapon Systems Support (NAVSUP, WSS) that calculates allowance levels for a pre-determined operational availability (A_o) level of systems or equipment and relates it to the sparing cost of aviation parts. NAVARM selects inventory levels for all parts (assemblies and subassemblies) in each weapon system (WS), from a set of allowance candidates. The selection seeks to minimize cost while ensuring the expected A_o is above the target threshold. The WS's program resource sponsor establishes the A_o level in accordance with Operational Navy Instruction (OPNAVINST) 4442.5A. In order for NAVARM to perform and give spares level recommendations to reach the target A_o, it makes several assumptions. Some of those assumptions are made for mathematical tractability of the problem, while others are based on operational considerations. A natural question arises about the importance to understand the sensitivity of NAVARM to violations of its assumptions.

A. PURPOSE

The intent of this thesis is to concentrate on some of the assumptions used by typical RBS models (Sherbrooke, 2004) like NAVARM to recommend inventory levels. Specifically, we perform a sensitivity analysis on one primary assumption: NAVARM's distribution of the mean time between failure (MTBF) for any repair part is assumed to be exponential.

We use the Readiness-Based Sparing Simulation (RBSIM) developed by Wray (2017) for the analysis. The simulation allows us to test the difference between NAVARM's estimated A_o and simulated A_o for recommended inventory levels as input parameters are varied. To test the distribution assumption, we modify the MTBF, changing its current exponential distribution into either Weibull, gamma, or log-normal distributions in RBSIM and compare the A_o results to the theoretical results NAVARM predicts.

B. BACKGROUND

RBS models derive from early supply-oriented optimization models following the end of World War II. The goal of these models is to minimize backorders for a given WS by recommending future spare parts levels. The earliest models focus on single items with a single inventory control point (Galliher et al., 1959). Soon after, the military services start to develop their own versions of these models to meet customer demand.

By the mid-1960s, the development of models that are able to calculate stock allowances from the assembly to the subassembly level (known as the multi-indentured model) for multiple items at a single site begins. Multi-indentured implies that a piece of equipment can have multiple sub-components that can be repaired or replaced. Muckstadt (1973) extends the capability of the single site model to account for the effects of subassembly backorders and how they affect higher indentured assemblies, and ultimately the availability of the WS. This model would evolve to incorporate the multisite hierarchal model (known as the multi-echelon model) into what is now known as the RBS model.

Multi-echelon implies that the repair of the failed part can occur beyond the organizational level at either the depot level or further up the supply chain at the manufacturer level. Multi-echelon also implies those locations where staging of ready for issue material to best support the expected A_o goal at minimum cost.

Today, the Navy uses NAVARM to plan spares order recommendation levels that ensure readiness thresholds are met in the most economical way. The Navy relies on NAVARM to produce the allowances for all aircraft carriers, 9 amphibious assault ships, 31 naval air stations, and 17 aviation platform packages for all of the Marine Aviation Logistics Squadrons (MALS). According to NAVSUP WSS's budget office, N8, the Navy spends on average $572M annually on these spares buys. NAVARM has the ability to process multiple WSs per project run, ranging from 1 (single WS) up to 24 (shore site consolidated allowance list). In terms of total aircraft numbers, NAVARM ranges from 1 to 352 in an individual project run.

C. THESIS STRUCTURE

The following is a description of the subsequent chapters:

- Chapter II discusses applicable literature related to the VARI-METRIC model that RBS uses in both civilian and military aviation, and the mathematics behind the demand distribution assumption.

- Chapter III reviews the input data as well as describe the RBSIM model event process and assumptions.

- Chapter IV analyzes the output data of RBSIM and compares it to NAVARM.

- Chapter V summarizes the conclusions and provides recommendations for future research on this topic.

THIS PAGE INTENTIONALLY LEFT BLANK

II. LITERATURE REVIEW

This chapter opens with a general description of the VARI-METRIC model. The Navy's RBS model, one instance of the VARI-METRIC model, provides the Navy with a functional allowance planning model. A discussion of the VARI-METRIC model is warranted for the context it provides later in this thesis. Next, we discuss why and how the VARI-METRIC model has been adopted for commercial and military aviation use. Lastly, we discuss the mathematical reasoning behind the MTBF distribution and the selection of alternatives. The selection of distribution alternatives is based on previous research of the VARI-METRIC model by Sherbrooke and is documented in his book (Sherbrooke, 2004, pp. 101–125).

A. VARI-METRIC MODEL

The VARI-METRIC model is first developed by Slay (1984). It improves upon the original Multi-Echelon Technique for Recoverable Item Control (METRIC) model that is first developed by Sherbrooke (1968). The METRIC model calculates the optimal stock level for every item in the WS across multiple sites. Slay notes that if his model could improve prediction of backorders, a user could consequently improve forecasts for parts needed in inventory (Slay, 1984). Sherbrooke extends Slay's work in his 1986 comparison of the VARI-METRIC model to that of the METRIC model (Sherbrooke, 1986, and 2004, pp. 101–125). The VARI-METRIC model is used for multi-indentured, multi-echelon systems. The indenture structure can be viewed as a tree, where the base of the tree is often referred to as the line replaceable unit (LRU) and the branches are known as shop replaceable units (SRU). Figure 1 depicts the basic LRU and SRU relationship.

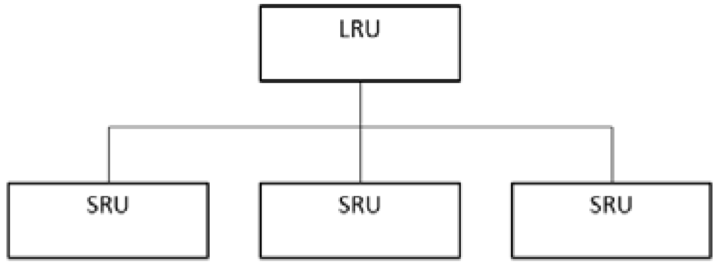

Figure 1. LRU and SRU Indenture Structure. Source: Sherbrooke (2004).

Sherbrooke uses the LRU and SRU relationship to explain how the VARI-METRIC model works (Sherbrooke, 2004, p. 102). When an LRU fails, it is removed from the system and sent to a depot (for repair) or purchase. If no spare parts are available, this action creates a backorder against the LRU. If the LRU is immediately repaired, the backorder is filled by the same LRU. Additionally, if an SRU is needed to repair the LRU, a backorder is also created against the SRU. Essentially, from the point of view of "pipeline" demand, two backorders have been created (one for the LRU and one for the SRU).

Repeating this process multiple times allows us to construct a theoretical distribution for LRU and SRU orders. This distribution's mean demand and variance at the depot level can be calculated for both the LRU and SRU. If we have reasonable estimates of delay time of repair or purchase for the LRU and SRU, we may estimate expected backorders (EBOs). Since maximizing A_o approximates to minimizing EBOs at a WS level, this allows us to build an approximate allowance table of parts needed at the depot to minimize cost for a certain level of A_o.

For repairable parts, Figure 2 depicts LRU and SRU relationship as LRU failure moves through inspection and repair stages.

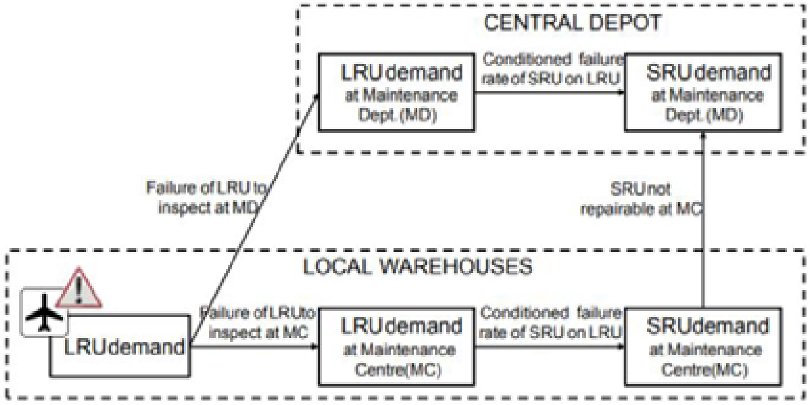

Figure 2. Event Diagram of the LRU and SRU Repair Process. Source: Constantino et al. (2013).

The VARI-METRIC model makes the following assumptions:

- One and only one SRU fails simultaneously.
- All SRUs that can be repaired are repaired at the depot.
- A part’s demand follows a Poisson distribution.
- The pipeline to calculate the EBOs follows a negative binomial distribution.

B. VARI-METRIC MODEL USE IN AVIATION

The aviation industry has experienced a constant increase in the volume of commercial flights since its inception. Consequently, the increasing volume of aircraft flights has brought about an increasing failure of aircraft equipment. However, both storage space and budget constraints limit the amount of spare parts the industry can have readily available in inventory. Sherbrooke’s multi-echelon, multi-indentured model is used in the aviation industry to help identify the best combination of spare parts allowances to keep on the shelf.

In 2013, Northwestern Polytechnical University School of Aeronautics in China published a report discussing the use of a multi-echelon inventory allocation model with a finite repair capacity that aims at optimizing aircraft spare parts for civil aircraft (Li et al., 2013). They focus on optimizing the maintenance resources for civilian aircraft. The research team utilize a M/M/c queuing model to study the effects of repair time on the maintenance cycle of civil aircraft. They conclude that the Sherbrooke model approach to aviation maintenance can reduce overall maintenance cost as well as improve the support structure of civil aircraft operations in comparison to the traditional spare parts inventory allocation model currently in use by the Chinese civil aviation industry.

At the same time, Costantino et al. (2013) analyze the VARI-METRIC model for spare part allocation to achieve an operational target given a budget constraint. In order to overcome that both the budget constraint and repair centers that have different capabilities, they develop the model to minimize the system-wide, expected backorder levels with a solving algorithm based on marginal analysis. The model determines the stock levels at each warehouse as well as the center depot. The model design they develop is for military aviation use. However, they argue that a commercial buyer can run the same model to calculate the best economic level of parts acquisition at the beginning stages of logistic support.

Both Li et al. (2017) and Constantino et al. (2013) determine that a VARI-METRIC model produce optimal results for minimizing aircraft downtime. More and more industries, such as ship-repair and maintenance facilities, are implementing their own version of Sherbrooke's model. Most employ the model with variations and additions on the assumptions.

C. NAVY ADAPTATION OF RBS FOR AVIATION

Today's Navy RBS model NAVARM connects the investment in spare parts to WS readiness. A WS down for a lack of parts is referred to in the military as non-mission capable supply (NMCS). Naval aviation consolidated allowance list (AVCAL) and shore consolidated allowance list (SHORECAL) allowancing uses historical demand to establish a base failure rate per item. This failure rate is applied to the aircraft population and projected flying hours to determine to failure rate at a specific site. In turn, NAVSUP outfits embarked airwings and shore facilities with necessary critical parts to maintain their readiness levels.

In accordance with the Chief of Naval Operations (CNO) instruction OPNAVINST 4442.5A, “RBS models should directly compute both range and depth for all echelons of supply” (OPNAVIST 4442.5A, p. 5). Range implies the number of replaceable or repairable parts needed while depth implies the quantity of each part. Furthermore, the RBS model is designed as an item-indentured structure for use by NAVSUP. Top, “parent” level, inventory items are identified as those whose next higher assembly is the WS itself, and are commonly referred to (by NAVSUP) as weapon replaceable assemblies (WRAs), instead of LRUs. Lower level items below WRAs are known as shop repairable assemblies (SRAs) instead of SRUs. These “child” parts are used to repair the parent parts.

The U.S. Navy began to invest into the RBS model during the mid-1980s. Their goal was to adapt the expected backorder models to forecast the level of spare parts that are needed to minimize backorders. In 1987 the CNO, Admiral Trost, directed NAVSUP to implement the use of RBS. He also directed aviation supply to embrace the idea (Naval Inventory Control Point, 2008, p. 4). Soon after, the Operational Analysis Department for NAVSUP in Mechanicsburg, PA, developed and implemented an RBS model to create AVCALs for aviation platforms. Strauch (1986) developed the RBS model known as the Aviation Readiness Requirements Oriented to Weapons Replaceable Assemblies (ARROWS).

The ARROWS model is a site-level stockage model that optimizes aviation parts in the multi-indenture structure. It works to reach a given part A_o constraint while minimalizing cost. The model does allow for multiple WSs at a single site and it considers the impact of any parts that share commonality among the WSs on the overall readiness of the system. ARROWS optimizes a single WS before moving sequentially to the remaining WSs in the optimization process. It does consider previously set stock levels as it moves from one WS to the next. The initial testing of ARROWS model utilizes data collected from the 1986 deployment of the USS ENTERPRISE. Strauch compares readiness rates of both the F-14 and the SH-60 aircrafts that have been reported during the deployment and shows that ARROWS calculates the rates to within 10% of the actual observations (Strauch, 1986). NAVSUP also uses a commercial model called the Service Planning Optimization. In 2016, NAVSUP asked NPS to develop NAVARM, with the capability to improve the heuristic VARI-METRIC logic and incorporate persistence of legacy solutions (aka "churn" control). The "churn" control minimizes the change between the new candidate solution and a previous solution. It acknowledges the cost of inventory recapitalization, which is not captured in RBS models. For example, if the legacy solution for item A is to stock a quantity of 2, at a cost of $100, but the new candidate solution is to stock 0 of item A and 1 of item B at $100, the RBS solution would have the same value. However, in reality the execution cost of recapitalizing the inventory is an additional $100 in order to buy item B.

NAVARM solves single-echelon, multi-indentured scenarios, optimizing part allowance levels for naval aviation WSs. It embeds a heuristic algorithm in order to recommend spare parts reorder points. NAVARM assumes an (S-1, S) inventory model at the retail level. S represents the maximum allowance stock level at an individual site that is determined by NAVARM. S-1 represents the reorder point where the inventory decreases by one. NAVARM's underlying theory of calculating EBO for an item follows Sherbrooke's VARI-METRIC model. This means the model utilizes the Poisson distribution for the overall number of failures for a given system, but the sub-components reflect a Negative Binomial distribution for the failure rate (Sherbrooke, 2004, pp. 101–125).

NAVSUP operates with a more complex indentured structure than those shown in academic examples such as in Sherbrooke (2004, pp. 101–125). The WS is composed of one or several LRUs. The indentured levels of the SRUs stretch further down than the typical SRU components. The LRU and SRU components also can spread across multiple WSs. Sherbrooke admits that this type of structure does "complicate the computer programs substantially … [albeit] the basic logic is the same" (Sherbrooke, 2004, p. 114). The diagram in Figure 3 depicts three WSs, each one with a single LRU (parts "A," "J" and "Q," respectively), and the SRU relationship. Part of the complexity lies in the common parts and "chain of influence" depicted modeled (i.e., cannot be replaced) in "WS3." The term "chain of influence" is used for those parts that are in one WS that have an influence in another WS. In Figure 3, "G" influences the A_o of "WS3" through their shared commonality of "L" even though "G" is not indentured to "WS3."

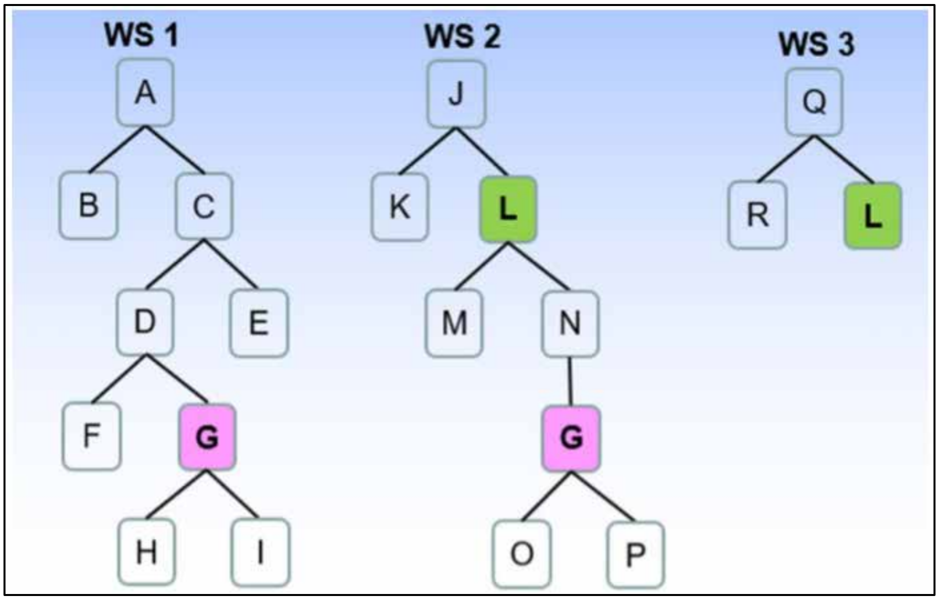

Figure 3. The Chain of Influence in a Multi-Indenture Part Structure. Source: Salmeron (2016).

NAVARM, like Sherbrooke's VARI-METRIC model, makes certain assumptions in order to optimize properly. These key assumptions are:

- NAVSUP's supplied formula for estimating the average WS readiness levels are based on EBO and supply system demand inputs which gives an accurate estimate of expected availability.
- Poisson is an adequate distribution for a part's failure disregarding its subpart failures. Negative Binomial distribution is an adequate distribution to compound the effect of subpart failures into the parent part failures.
- Sherbrooke's VARI-METRIC model framework approximates EBOs correctly.
- Partial mission capable WSs are not counted as being available. WS non-availability is a result of all parts failure.
- Cannibalization of parts from one WS to another does not occur.

D. DISTRIBUTION ALTERNATIVES

For the purpose of this thesis the terms MTBF and demand for parts are interchangeable. The naval supply part ordering system does not register a part failure but instead registers a demand order. However, the demand is often because of a failure. Therefore, MTBF and demand represent the same process.

In early VARI-METRIC models, the demand for parts assumes Poisson values with a mean estimated by a Bayesian procedure. However, a problem arises when the mean of the demand becomes non-stationary. This creates a mean to variance ratio (MTVR) greater than one. Sherbrooke discusses how this becomes problematic while using Poisson as the distribution estimate when trying to estimate parts especially with low demand (Sherbrooke, 2004, p. 89). He suggests that gamma, Weibull, or log-normal distribution may provide a better fit.

An article published by the Department of Mechanical and Aerospace Engineering at Sapienza University of Rome analyzes the performance of both Poisson and Weibull as the demand distribution (Patriarca et al., 2019). The intent of the article is to test the multi-indentured, multi-echelon model's demand distribution using the discrete Weibull distribution. The research team design a simulation to calculate backorders by utilizing a demand data set. The outcome of their simulation shows that traditional models based on the Poisson distribution do not necessarily reflect the best framework for some demand patterns. They note that the Weibull distribution performs better than Poisson in estimating backorders.

The above-mentioned paper utilizes Weibull as the alternate distribution for their model. However, Sherbrooke suggests that the gamma distribution may be an easier alternative (Sherbrooke, 2004, p. 89). Sherbrooke states "The demand process for some parts are not random, but results from wear and tear…the probability distribution of time to the next demand does not decrease uniformly like the exponential. Instead there is a peak value to the right of the origin as in distributions such as gamma, Weibull, or log-normal." (Sherbrooke, 2004, p. 89). The mean and MTVR can be specified to compute the parameters of gamma. The Weibull distribution's parameters are determined by solving two nonlinear equations. The Weibull distribution has an advantage over gamma: If the time until next failure or demand is determined probabilistically, the Weibull distribution is better suited to sample. This thesis utilizes both distributions as well as the log-normal distribution in the simulation in order to explore Sherbrooke’s theory of alternatives.

THIS PAGE INTENTIONALLY LEFT BLANK

III. DATA AND SIMULATION REVIEW

A major issue with the data for this study is that the Navy's method of collecting does not accurately reflect the mean time between failure for individual components. In order to track the data correctly, components would have to be serialized and monitored from the initial installation into the WS. Currently, the Navy only records part failures when the part is ordered in the supply system. NAVSUP WSS keeps track and records these orders in order to create the demand history of the parts. Therefore, the MTBF calculation is a result of a part's demand in the supply system and not the actual operating lifespan.

Next, we examine the simulation model of this thesis, RBSIM, developed by John Wray (2017). The aim is not to discuss the internal coding of the simulation, but to provide an overall understating of its operation and how it has been modified to use alternative distribution calculations.

A. DATA COLLECTION AND MANIPULATION

NAVSUP WSS populates a Microsoft Access file (known as the candidate file) with pertinent part requisitions information gathered from the supply system in order to capture demand. The part data NAVSUP WSS collects become the initial input for the aircraft carrier's AVCAL. NAVARM uses the candidate file as the data input, along with a few other algorithm specific inputs, and populates an output allowance table and is used as the input data for the simulation of this thesis. This thesis utilizes NAVSUP WSS's candidate file for the USS Carl Vinson (CVN 70).

To fit new distributions to the collected data, this research adds a few parameters that do not exist in the original data: First, we designate a demand distribution type: If the demand of the part is less than five units over the 24-month period of data collection, we assign an alternate distribution; if the demand is over five units, we retain the default exponential distribution. This thesis simulates the Weibull, gamma, and log-normal alternative distributions as well as the exponential default distribution. Second, we specify the MTVR. Sherbrooke takes note of the importance of this ratio's use for analysis when he states “Our interest is not the time until the next demand, but something related to it – the mean and the mean-to-variance ratio for the number of demands over the pipeline.” (Sherbrooke, 2004, p. 90). We note that a MTVR of exactly 1.0 corresponds to an exponential distribution for demand inter-arrival times.

For the alternate distribution we assign a MTVR value that is not equal to 1.0. We accomplish this by running the simulation with an alternate distribution at a MTVR of 1.5 then run the same simulation with a MTVR of 0.5. Table 1 shows the type of distribution with the MTVR assignment. The ratios of 1.5 and 0.5 are arbitrarily chosen as numbers that are either greater or less than 1.0. We begin by assigning all parts in the database with the same alternate distribution. Once all of the alternate distributions complete the simulation, we incorporate the alternate distributions with the original exponential distribution in the database. We determine the distribution of the part based on the demand level. We classify low demand as having five or less requisitions in the candidate file we obtain from NAVSUP WSS. We assume a MTVR of 1.5 and 0.5 for low demand items. These mixed distributions contain both the alternate and default exponential distribution.

Table 1. MTVR Selection for Sequential Simulations on the Three Different Distributions Criteria.

MTVR Selection for Sequential Distribution Simulations Criteria		
Distribution Criteria	MTVR Selection	
	First Simulation for Each Distribution	Second Simulation for Each Distribution
Default	1.0	None
100% Alternative	1.5	0.5
Mixed Alternative	1.5	0.5

NAVARM calculates the EBO and populates the NAVARM output allowance table with each SRU parts allowance that becomes the basis for the AVCAL. The allowance output table is also the input data for the RBSIM simulation model that the research team uses to test the underlying MTBF distributional assumption. RBSIM uses the NAVARM allowances and specified distributions and calculates EBO and A_o by simulating part failures and replacements.

B. SIMULATION OVERVIEW

The RBSIM simulation was developed by Wray (2017) to help verify NAVARM's outputs as well as provide additional data outputs for further study by analysts and decision makers. The model uses NAVARM's allowances for all parts, and simulates failures at the individual SRU level. Then, it aggregates the parts' chain to the parent LRU in order to estimate the WSs EBO and A_o. We compare RBSIM's A_o results to NAVARM's estimates. During the simulation, RBSIM characterizes parts by:

- Status (whether the part is operational or not due to maintenance and/or supply)
- Planned failure times
- Physical position (if in use where is the part installed)

In addition, each WS is characterized by:

- Aircraft type
- Operational status
- List of the part positions within the WS

Lastly, each part position is characterized by:

- The WS, if it is currently in use
- Expected failure times parameters
- Parameters for individually repaired parts to return to inventory

RBSIM calculates expected completion times for repair and subsequent return to inventory based on the type of failure. RBSIM sums the different type of failure rates to form a combined failure rate. A random number draw compares the ratio of repairable and non-repairable failures in order to assign the type of failure when one occurs.

Wray's original RBSIM generates a stochastic failure time of the individual parts based on the exponential distribution. We have extended RBSIM so it can use either Weibull, gamma, or log-normal-distributed time between failures, where the inputs to these distributions are now based on mean and variance as determined by the MTVR.

The failure rates are also specific to a part's position on a particular WS. For example, a circuit card that is in an F/A-18 may have a different failure rate than the same circuit card in an SH-60. The circuit card may even be in multiple LRUs across multiple aircraft platforms and experience different failure rates for each SRU based on the part's location within the WS. When a part fails, RBSIM immediately requisitions a new part and removes the failed part from circulation. It keeps the part out of the rotational pool until the part-specific completion time of repair.

The basic steps of RBSIM's core logic are:

- Reading in the data from NAVARM's allowance output.
- Assigning a first-time failure rate of each part based on the input distribution, and assigning parts to fill each WS.

When a failure occurs, RBSIM takes action by:

- Assigning the return to service time.
- If inventory is available, decrease the inventory level by amount of the part failure and place the WS in a "down" status for the duration of the specific mean time to repair (MTTR). If the inventory is not available, RBSIM adds the WS to a first-in first-out (FIFO) queue for the part.

Lastly, when the part is repaired and available for issue:

- RBSIM uses it to repair the first WS in the FIFO queue; and,
- If no WS is awaiting repair, the repaired part returns to inventory.

RBSIM uses summary levels of data produced from recorded flight hours of the aircraft, parts failure intervals, repair times, and parts transportation lead times as deterministic values. These expected values are used in lieu of an actual repair process cycle at an intermediate or depot level repair site because the actual cycle has minimal effect on the metrics of interest. The simulation also ensures that the simulation run-time is within acceptable parameters by tightly scoping the factors that are taken into account in the simulation without having to sacrifice much fidelity for the metrics of interest.

RBSIM only calculates metrics that are related to readiness in order to reduce run time and code complexity. RBSIM's output is straightforward, and consists of the following metrics:

- The mean of the backorders by part type
- The mean of the on-hand inventory per part type
- The fill rate of each part type
- The average A_o of each WS used in the simulation
- The percent of time the A_o was at or above the specified availability goal per WS type

C. SIMULATION EVENT GRAPH OVERVIEW

In this section, we explain the RBSIM originally developed by Wray (2017). Figure 4 shows a simplified event graph. It describes the overall process of the part failures and ensuing repairs. The event graph provides a broad level of understanding of how the parts flow through the simulation. RBSIM utilizes the Java Simkit Library to implement the simulation calculations (Buss, 2019). Simkit converts the event graph in Figure 4 into working computer code to support the simulation. The open source UCanAccess (2017) library also interacts with NAVSUP's Microsoft Access database in support of RBSIM.

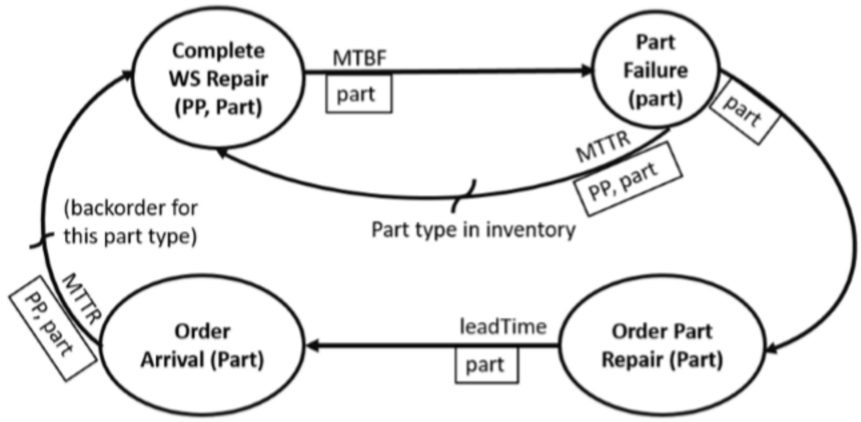

Figure 4. RBSIM Event Graph. Source: Wray (2017).

The simulation begins at time zero where it initializes all state variables for the system. It assigns each part position a part on each WS in order to have each WS begin in an operational status. Known as the Run event (not depicted in Figure 4), the initialization process also schedules the *Part Failure* event for each part at a specific time using the data that is available for each of the part's position. The Run event then calculates the next expected time to failure for the part based on the part's position then provides the calculated parameter for the random number draw. The Run event completes when every part in the WS has one scheduled failure.

The *Part Failure* event simulates the failure after it receives the failed part parameter from the initialization. This event sets both the part and the parent WS statuses to non-operational and schedules an *Order Part* event. If the inventory has the part available, the *Part Failure* event schedules a *Complete WS Repair* event to commence by using the MTTR associated with the specific WS. If the inventory does not have a ready-for-issue part, the part type is added to a FIFO queue by part position.

The *Order Part Repair* simulates acquiring parts from the supply system. The supply system turns the part in for repair and receives a part from inventory or from the repair cycle. This event also calculates an expected shipping lead time and subsequently schedules the *Order Arrival* event for each part that is in the cycle. A random number draw determines the shipping time and compares it to the repairable parts ratio to determine if the shipment time should be calculated for repair at the given single-site or a resupply from the depot.

The *Order Arrival* is the event where the site supply system receives the ready-for-issue part. The corresponding inventory increases by one upon receipt of the part. If there is an outstanding order against the part in the FIFO queue, an MTTR delay is applied and the part is sent to the *Complete WS Repair* event.

The *Complete WS Repair* event occurs when part installation occurs. It also generates a new time to failure using a random number generator based on the part's input distribution. Lastly, the event checks all part positions in the WS to make sure all parts have been assigned. If all part positions have corresponding functioning parts, the WS status changes to "up" meaning the WS is operationally available.

D. SIMULATION ASSUMPTIONS OVERVIEW

In order for RBSIM to run in a reasonable time, certain assumptions are taken in the implementation of the simulation. Some of these assumptions could impact results but are made to keep run-times reasonable. The RBSIM principle assumptions are:

- Failure rates are accurately represented by a specific distribution. This is the primary assumption this thesis is testing. As stated earlier, Weibull, gamma, and log-normal will replace the exponential distribution for some parts.

- Failures are independent from one another. Scheduled failure times are on a continuous timeline and there are zero-part dependencies within the simulation. Simultaneous failures will not occur in the simulation although this sometimes occurs in the real world.

- Simulated failures will continue to occur even though the WS is listed as non-operational. This is to ensure that scheduled failures continue to occur in the simulation and that the expected failure rate is upheld.

- SRA part failure times do not reset when replacing the parent WRA. The assumption is that when the part repair occurs, it does not affect the dependability of the other SRAs in the WRA. The collected data does not define how often parts repair is completed at a separate installation and what happens to the SRAs when an inventory restock is necessary for the parent WRA. The assumption does lead to a conservative estimate level of readiness; however, the degree of the impact on A_o is not known at this time.

- Demands are set as FIFO. No priority parts demand from the fleet has been given to part orders. Also, priority has not been given to WS's that are below the specified A_o goal.

- Lateral supply support is not allowed. Sites that have high inventory cannot fill requisitions from sites that have low inventory.

- The practice of removing parts from one WS to another WS in order to return non-operational WS back to an "up" status does occur in the real world, and is known as cannibalization. However, NAVARM achieves the desired readiness levels the user chooses without utilizing this practice so therefore RBSIM does not allow cannibalization of parts.

- Repair times are independent from one another. RBSIM does not simulate a backlog of the repair pipeline involving multiple parts of the same type that fail simultaneously.

- Demand rates are stationary for the simulation horizon.

In future analysis using RBSIM, these assumptions can be replaced to better represent real world operations. This thesis tests the impact of the first assumption on A_o.

IV. ANALYSIS

In this chapter we discuss the results produced by RBSIM, with the aim of analyzing whether A_o differences exist across WSs, and/or if alternate distributions and MTVRs make a difference to the achieved A_o. Before doing so, it is important to note that we observe a modest discrepancy between the A_o reported by NAVARM and the A_o estimated by RBSIM using the exponential distribution: RBSIM produces consistently higher A_o estimates than NAVARM. This difference appears to increase as the target A_o decreases. Since RBSIM is not an official NAVSUP tool, and has been tested only by its developer, Wray (during his thesis research), we cannot guarantee its A_o estimates are accurate. Further examining RBSIM is outside of the scope of this research. Therefore, we will proceed with the assumption that even though RBSIM estimates may not be accurate, at least they are "similarly biased" for all runs we perform. For example, if demand distribution "A" produces an RBSIM-estimated A_o higher than "B," we assume that the difference in A_o is still a reasonable approximation.

To begin our analysis, we execute RBSIM for 30 replications for every setting (distribution, MTVR) and calculate the average and confidence interval. Table 2 is an example of the F/A-18E WS mean A_o and confidence intervals. In all cases the reported A_o is reasonably representative, however, for the sake of brevity we choose to omit the confidence interval tables for the remaining WSs. RBSIM also records what percentage of parts are of the alternate distribution when a distribution mix is used. As stated before, we repeat this process for each of the three alternative distributions. Each replication of RBSIM produces a single observation of A_o based on the percentage of time the given WS was operational. We average A_o over the 30 simulation replications. We continue by analyzing the results first by WS then by distribution. Finally, we perform a statistical examination of the data.

Table 2. F/A-18E RBSIM Distribution 30 Replications Mean A_o and 95% Confidence Intervals.

F/A-18E WS 30 Replications Mean Ao and 95% Confidence Intervals		
Distribution	**A_o Mean**	**95% Confidence Interval**
Exponential	87.98%	(87.87%, 88.09%)
Weibull (1.5)	74.01%	(73.77%, 74.24%)
Weibull (0.5)	80.43%	(80.26%, 80.49%)
Weibull Mix (0.5)	79.87%	(79.78%, 79.97%)
Gamma (1.5)	83.57%	(83.39%, 83.76%)
Gamma (0.5)	89.67%	(89.58%, 89.76%)
Gamma Mix (1.5)	74.36%	(74.15%, 74.58%)
Gamma Mix (0.5)	79.76%	(79.64%, 79.89%)
Log Normal (1.5)	86.67%	(86.55%, 86.80%)
Log Normal (0.5)	89.65%	(89.57%, 89.73%)
Log Normal Mix (1.5)	76.32%	(76.17%, 76.48%)
Log Normal Mix (0.5)	79.85%	(79.70%, 79.80%)

A. RESULTS BY WS

The WS analysis separates the WSs by their assigned mission and aircraft type aboard the carrier. For example, the F/A-18 WSs are together because they share the same type of aircraft although they are different models. The H-60 helicopters grouping is based on the same criteria as the F/A-18. Lastly, the E-2D and V-22 analyses are together because both aircraft serve in support roles on the carrier even though they are not the same type of aircraft.

Figures 5, 6, and 7 show the impact of each distribution and MTVR on the three F/A-18 aircraft design WSs. The red line indicates the target A_o. As a reminder, the "Mix" cases indicate that demand is assumed to: (a) follow the specified distribution and MTVR for parts with less than five units (over the 24-month period of data collection); and, (b) follow the original exponential (MTVR=1) distribution for the other parts.

What is most surprising is that, although all three WSs share the same integrated logistical support pipeline from the same manufacturer, all the estimated A_o are different for each distribution. This can possibly be attributed to the age of the WSs and mission roles each WS plays aboard the carrier. For instance, the EA-18G WS main mission is electronic warfare while the F/A-18 E/F WSs main missions are primarily air-to-air and air-to-ground combat. The F/A-18 E/F WSs are also older than the EA-18G WS and may not have as robust logistical support chain as the newer EA-18G.

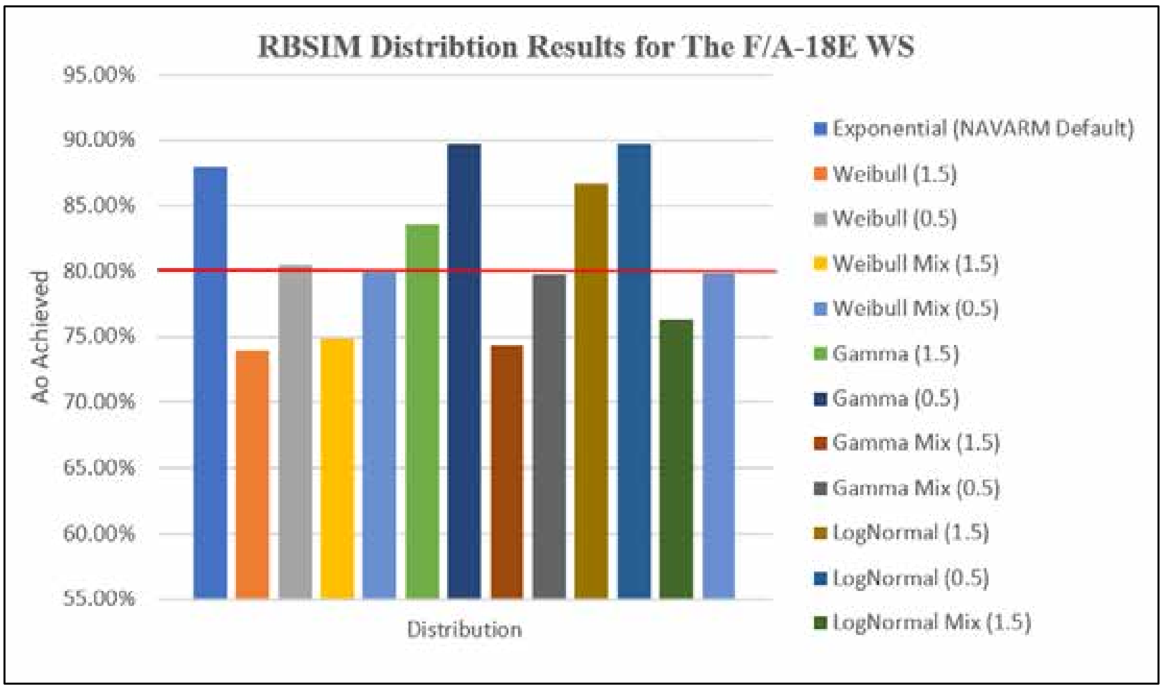

Figure 5. F/A-18E WS RBSIM A_o Results on All Distribution Combinations.

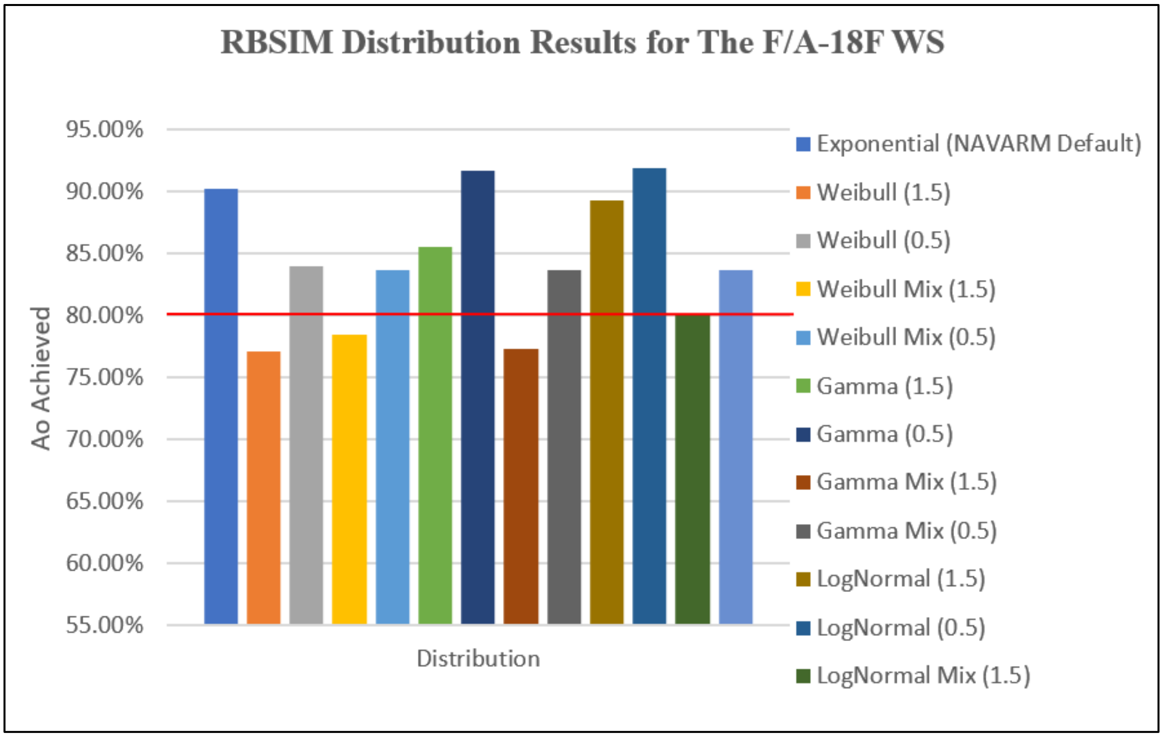

Figure 6. F/A-18F WS RBSIM A_o Results on All Distribution Combinations.

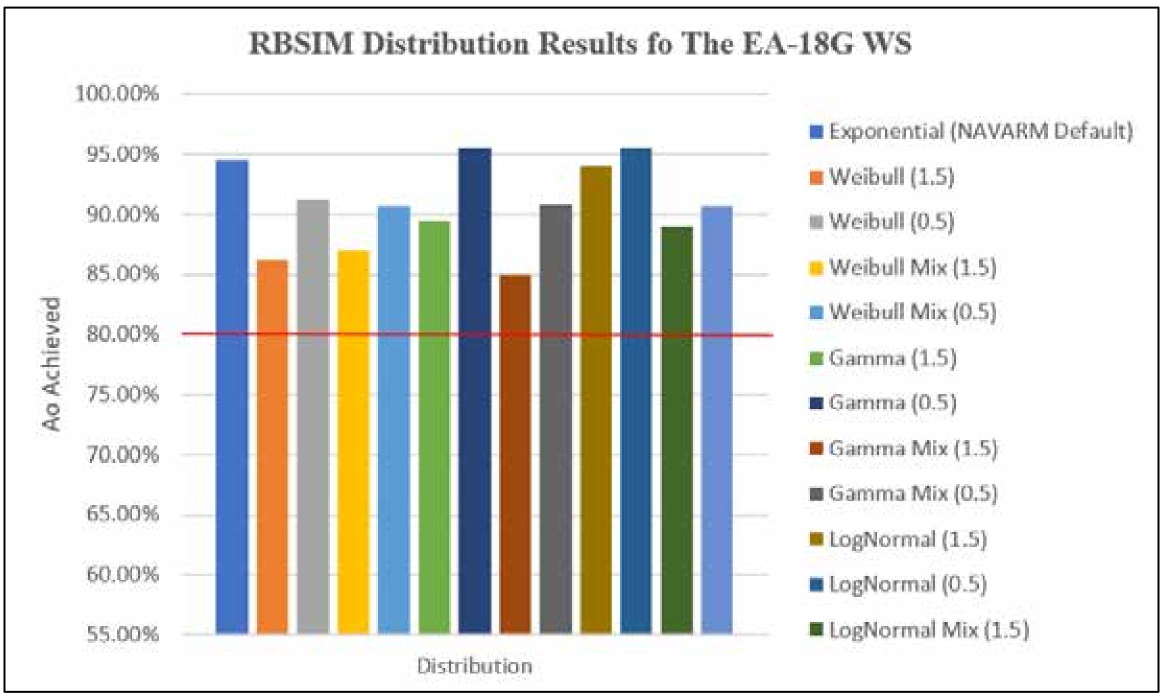

Figure 7. EA-18G WS RBSIM A_o Results on All Distribution Combinations.

It is interesting to note that F/A-18E WS achieves the lowest A_o results among the three aircraft models. The F/A-18E WS only achieves target A_o in 46% of the distributions (those distributions at or above the red line in Figure 5) compared to 76% for F/A-18F and 100% for the newer EA-18G WS. However, comparing the A_o graphs of all three together, it appears the F/A-18 E/F WSs closely resemble each other per distribution while the EA-18G does not resemble either of the previous two WSs even though they are all the same type of aircraft.

For the “mixed” case, 47% of the F/A-18E WS parts convert to an alternate distribution while 53% remain exponential. The F/A-18F WS has 47% of the parts converted and the EA-18G has 52% of all WS parts converted to the alternate distribution. Surprisingly, all three alternate distribution mixes have lower A_o results than the distributions where 100% of the demand was non-exponential regardless of MTVR. Since the default distribution in NAVARM is exponential, the alternate distributions that are combined with exponential should yield a higher A_o than their 100% alternate counterparts. This should be expected since NAVARM’s output allowance levels achieve the target A_o utilizing the exponential distribution. The gamma exponential mix distribution with the higher MTVR of 1.5 appears to achieve the lowest A_o among all the different mixed distribution. The Weibull mix with high MTVR achieves the second lowest A_o.

Overall, the F/A-18 E/F WSs perform as expected with the alternate distributions producing lower A_o than the original exponential distribution. The EA-18G WS achieves the highest A_o of all the WSs tested. The A_o never falls below 83%, well above the 80% requirement. This can possibly be attributed to: (a) the potential inaccuracy of RBSIM A_o estimates; and, (b) the normal wear and tear of parts on the EA-18G being not as pronounced as it is on the earlier models of the F/A-18 aircraft design. Therefore, the E-18G WS does not consume the amount of the allowance parts levels to have A_o below the target goal.

Figures 8 and 9 show the results of RBSIM on the H-60 WSs. Although both WSs are from the same type of aircraft and the S model immediately follows in design of the R model, they perform completely differently from one another. The H-60S WS achieves target A_o for 76% of simulation distributions which is shown in Figure 9 as those distributions that are above the red line; the H-60R WS does so in 30% of the distributions.

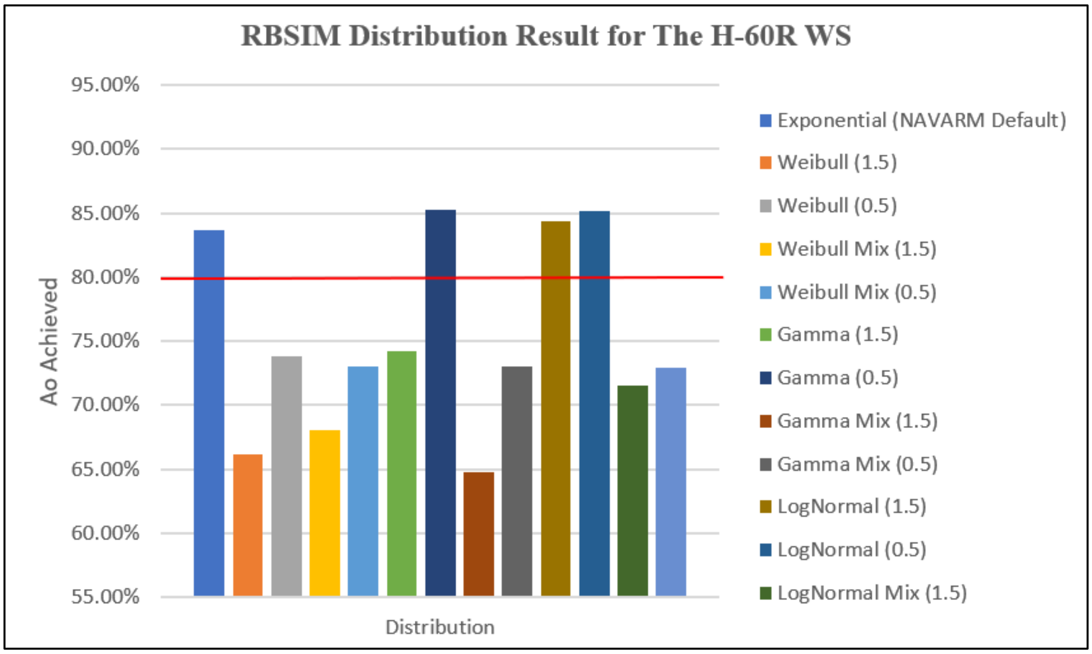

Figure 8. H-60R WS RBSIM A_o Results on All Distribution Combinations.

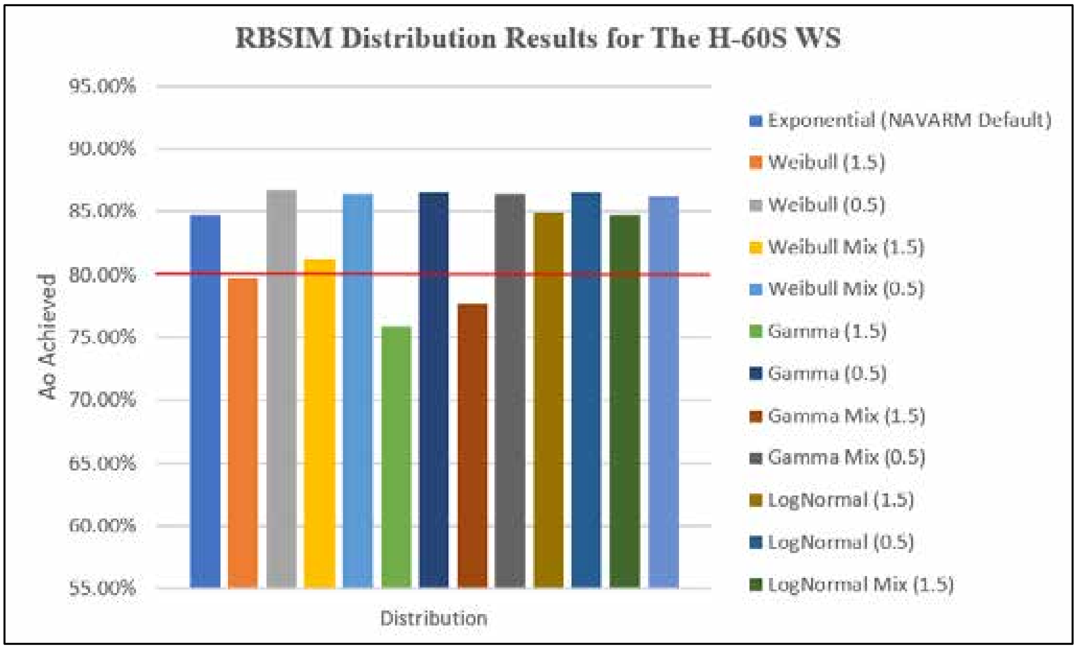

Figure 9. H-60S WS RBSIM A_o Results on All Distribution Combinations.

It appears the gamma distribution with a low MTVR of 0.5 achieves the highest A_o among all the distributions for both WSs. It even surpasses the default exponential distribution for both WSs by 2%. This is understandable since the MTVR of exponential is 1.0, so the variability is substantially less for MTVR of 0.5. The 100% alternate log-normal distribution with the MTVR of 1.5 achieves the second highest A_o for both WSs. The H-60R alternative distribution mixes achieve the lowest A_o among the distributions just like in the F/A-18's case. The H-60S WS alternative distribution mixes achieve a higher A_o than the 100% alternate distributions.

Both of the H-60 aircraft WSs have the same logistics support chain as well as share a lot of the same parts, yet perform completely differently in RBSIM. For instance, the difference in the parts change for the mixed alternate distributions is minor. The H-60R experiences a 49% change and the H-60S experiences a 47% change. Yet, the H-60S WS appears to be able to achieve a higher A_o across all of the distributions more often than its older platform model. This could possibly be explained by the age of the parts on the H-60R WS is becoming a factor and that wear and tear is more often than not the culprit of the additional demand.

Lastly, we analyze the support aircraft and how they perform with the alternative distributions. Figures 10 and 11 show the results of the how the alternative distributions affected the A_o of both WSs. The E-2D WS is the newest variant of the E-2 aircraft yet it achieves the lowest A_o among the seven WSs that are tested. The V-22 WS is the Navy's variant of the Marine Corp's V-22 combat support aircraft. Its performance is consistent with middle of the group of the WSs. Like the F/A-18 E and F variants, it appears the V-22's allowance levels built by NAVARM are robust enough to handle the majority of the alternate distributions.

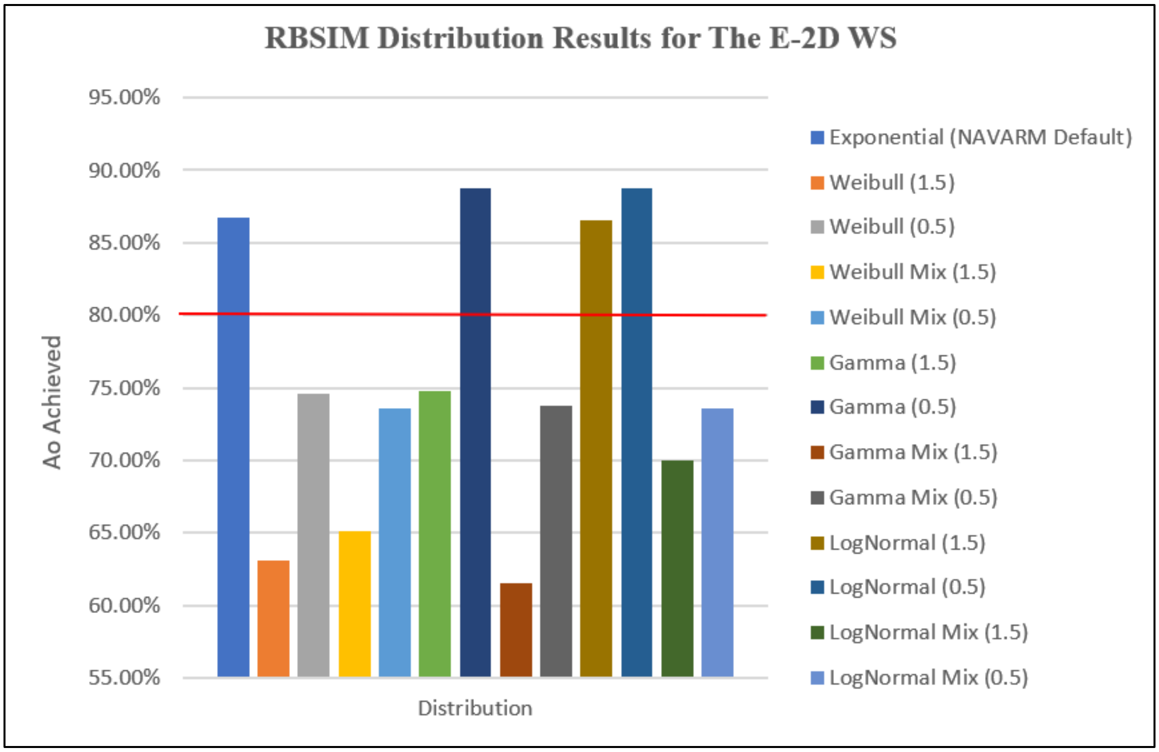

Figure 10. E-2D WS RBSIM A_o Results on All Distribution Combinations.

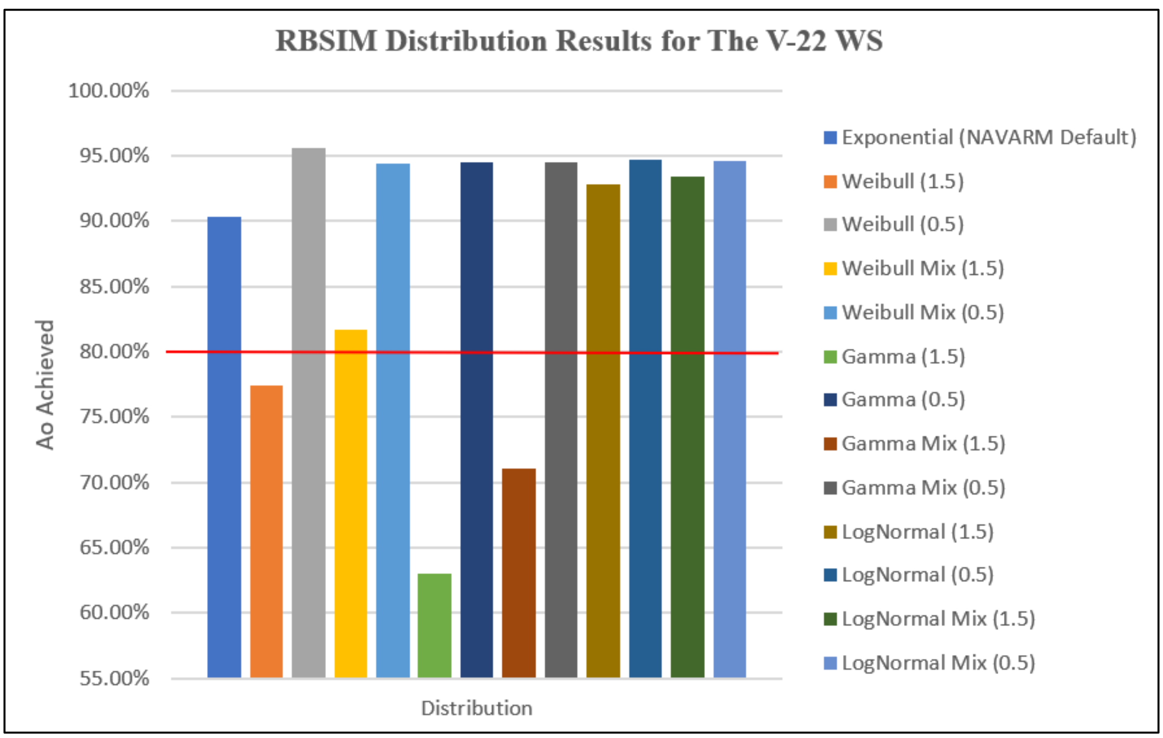

Figure 11. V-22 WS RBSIM A_o Results on All Distribution Combinations.

The 100% alternate log-normal distribution with an MTVR of 0.5 achieves the highest A_o for the E-2D at 89%. The Log-normal distribution A_o is higher than the default exponential distribution by 2%. For the V-22, the Weibull distribution with an MTVR of 0.5 achieves the highest A_o at 96%. It appears that in both WSs the gamma distribution achieves the lowest A_o. However, in the E-2D WS it is the gamma mix with the MTVR 1.5 that achieves the lowest A_o at only 62%. For the V-22 it is the 100% alternate gamma distribution with an MTVR of 1.5 that achieves the lowest with a 63% A_o.

Both WSs experience the majority of their unique parts change to reflect an alternate distribution. In particular, 61% of the E-2D WS parts change and 64% of the V-22 change to non-exponential for the distribution mixes. However, the E-2D is not consistent with the performance of the distributions. In particular, the Weibull mixes achieve a higher A_o than the 100% alternate Weibull distributions for both the higher and lower MTVRs. Yet the gamma and log-normal mixed distributions achieve lower A_o results than their 100% alternates. The V-22 mixed distributions achieves higher A_o than all of the complete alternates.

For all WS except the E-18G, using a high MTVR (1.5) results in lower estimated A_o than the target A_o. This suggests that, if actual failure time distributions do have a higher MTVR than 1.0, then the achieved A_o will be lower than projected. Without this analysis, we can only observe that the MTVR used to create allowances appears to make a difference in actual A_o achieved.

Lastly, we group the WSs and evaluate the A_o results based on the MTVR. Figures 12 and 13 shows the results. The graphs show that the 1.5 MTVR consistently achieves a lower A_o than that of the 0.5 MTVR. The mean A_o for the 1.5 MTVR is 79% while the mean for the 0.5 MTVR is 86%. Only three WSs achieve the target A_o with a 1.5 MTVR while all but two achieve the same target A_o with a 0.5 MTVR. It appears the MTVR makes a difference in achieving the target A_o for WSs.

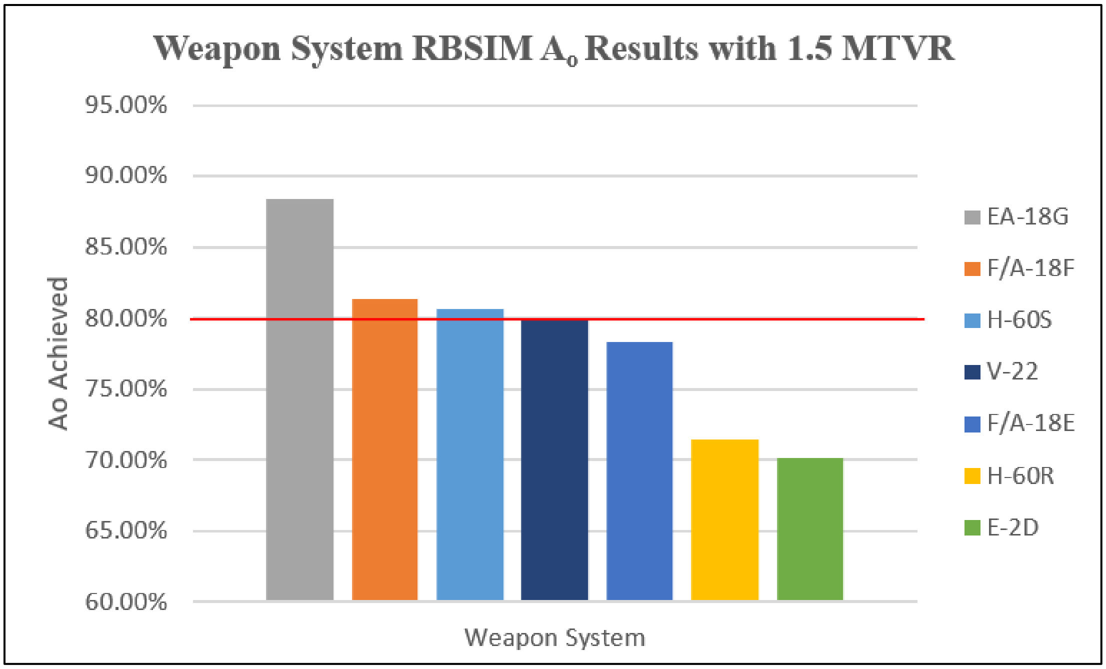

Figure 12. RBSIM Collective A_o Results for all Weapon Systems at 1.5 MTVR.

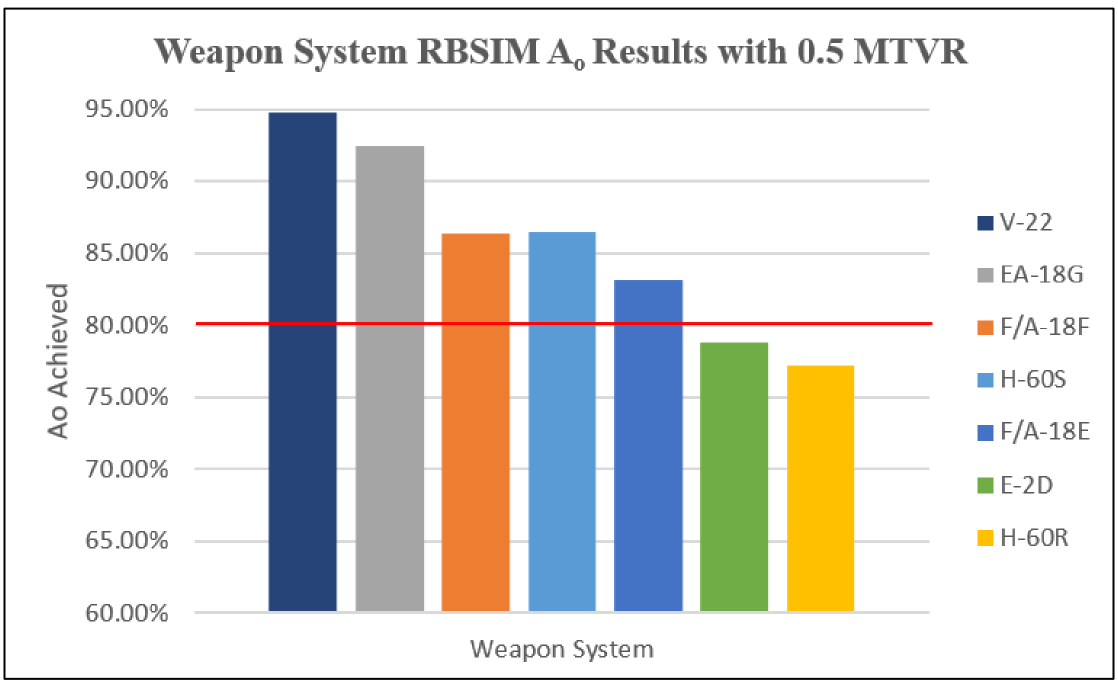

Figure 13. RBSIM Collective A_o Results for all Weapon Systems at 0.5 MTVR.

B. RESULTS BY DISTRIBUTION

In this section we discuss results of each distribution across all seven weapon systems. The goal is to understand the trends of each distribution by grouping the results across the WSs. Just as important is to see which distribution produces similar results as when using the exponential distribution. If an alternative distribution achieves the same A_o as exponential then the use of exponential distribution can be seen as a reasonable one. On the other hand, a discrepancy suggests that the exponential assumption may be suspect.

First and foremost is the performance of the exponential distribution in RBSIM. Figure 14 shows the results of RBSIM's use of NAVARM's allowance levels based on the exponential demand distribution. All of WSs achieve the target A_o. These levels are the basis achievements of the simulation that NAVARM's output allowance levels are designed to achieve. In other words, the exponential simulation closely resembles how NAVARM's allowance calculations are designed to perform in the real world.

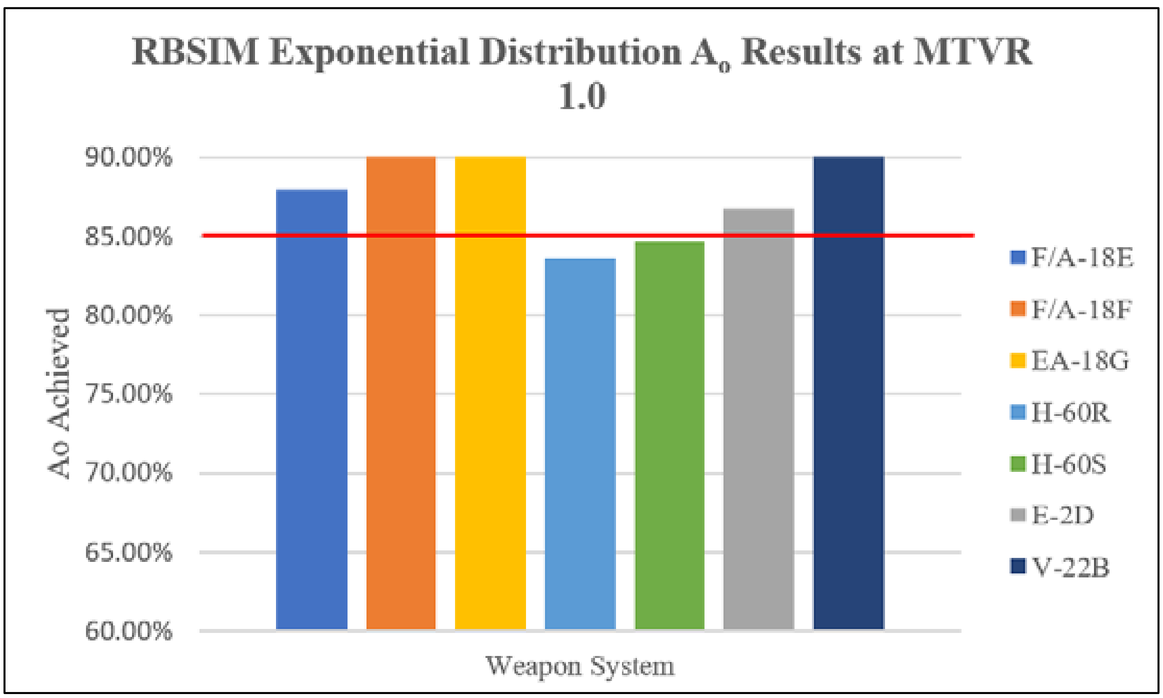

Figure 14. RBSIM Exponential Distribution A_o Results by WS.

Figure 15 shows the results of Weibull distribution (for 100% of the parts) with both MTVRs, as well as the mixed (Weibull and exponential) distributions. The 100% Weibull distribution appears to achieve the lowest A_o overall among the alternate distributions. Also apparent is that the Weibull distribution achieves a higher A_o with a lower MTVR. When all parts change to reflect a MTVR of 0.5, 57% of the WSs achieve target A_o. The Weibull and exponential parts demand distribution mix follows closely with 42% of the WSs achieving target A_o, with only the F/A-18E WS missing the goal by less than 1%.

When the MTVR increases to 1.5 the achieved A_o is much lower. Only one WS achieves target A_o when the entire demand distribution changes. The H-60S WS is within 1% of the target A_o. The mixed Weibull distribution with MTVR of 1.5 appears to produce higher results than its 1.5 MTVR counterpart, where 42% of the WSs are able to achieve target A_o.

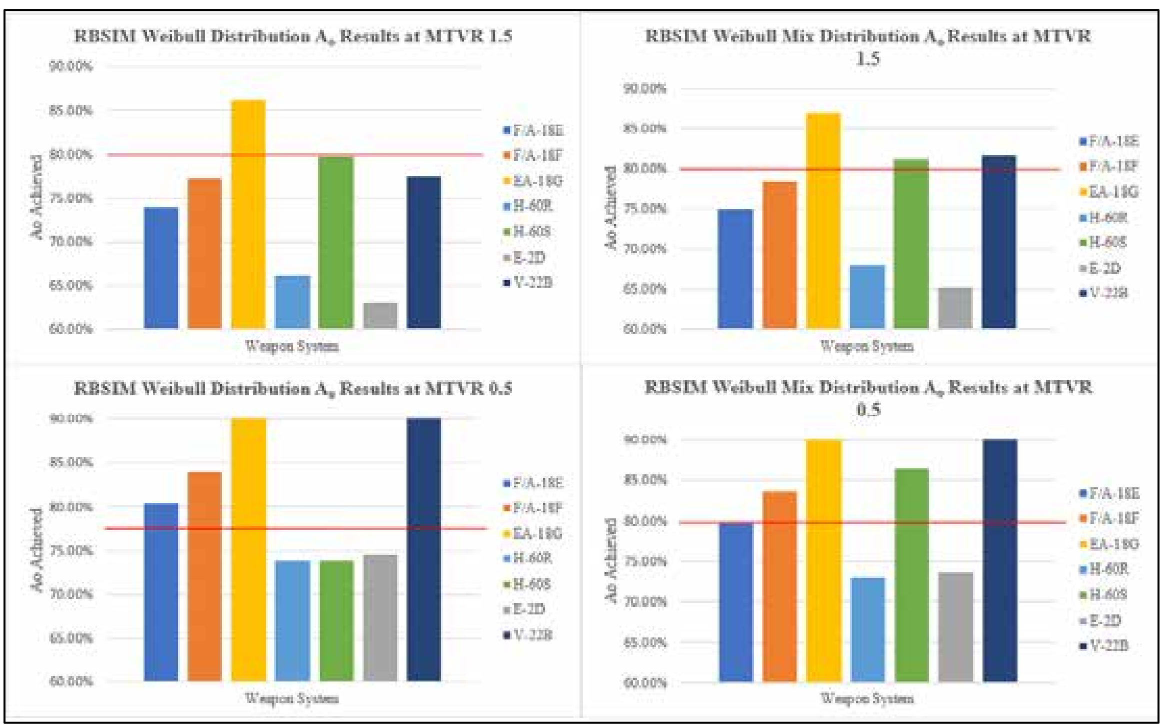

Figure 15. RBSIM Weibull Distribution A_o Results by WS.

The gamma distribution performs very differently depending on the MTVR (see Figure 16). A lower MTVR appears to produce the target A_o more often. Specifically, 85% of the WSs for the lower MTVR achieve the target A_o, with the F/A-18E reaching within 1% of the goal when the distribution is 100% gamma. The gamma mixture with the low MTVR produces the second highest results where 71% of the WSs achieve target A_o.

When applying a large MTVR to the gamma distribution the majority WSs do not meet target A_o. The difference in failure appears quite significant when the MTVR of 1.5 (where only 29% of the WSs achieve the target A_o) is compared to the performance of the 0.5 MTVR (where 71% do so). It is also quite surprising that the introduction of the default distribution decreases the A_o among the WSs. When the exponential distribution is mixed in with the gamma distribution for rate of failure, only 36% of the WSs achieve the target A_o while 64% achieve A_o for the100% alternate gamma distribution.

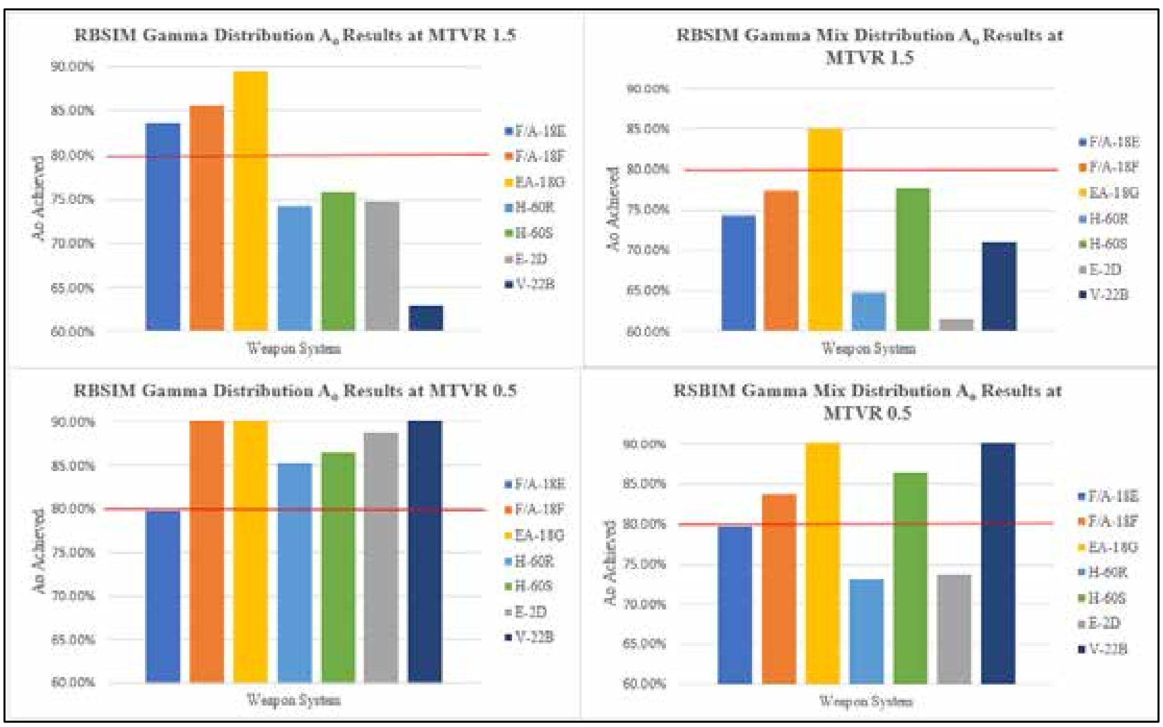

Figure 16. RBSIM Gamma Distribution A_o Results on WS.

Log-normal distribution appears to produce the highest A_o results of all three alternative distributions. Figure 17 shows the results of log-normal in RBSIM. Specifically, both log-normal distributions (where 100% of the part's distributions change) perform just as (well if not better) than the exponential distribution. All WSs are able to achieve the target A_o. The log-normal distribution with the lower MTVR is the highest performing distribution among all those that are being tested, including the default exponential demand distribution.

We have seen consistently that the mixed distributions do not achieve the same levels as the 100% alternate counterparts. However, both of the log-normal distributions perform just as well as the other two alternative mixed distributions. In all four log-normal simulations, the EA-18G and the V-22 produce the highest A_o results with both achieving 90%. Of the three alternate distributions, log-normal appears to achieve relatively the same results as the exponential distribution.

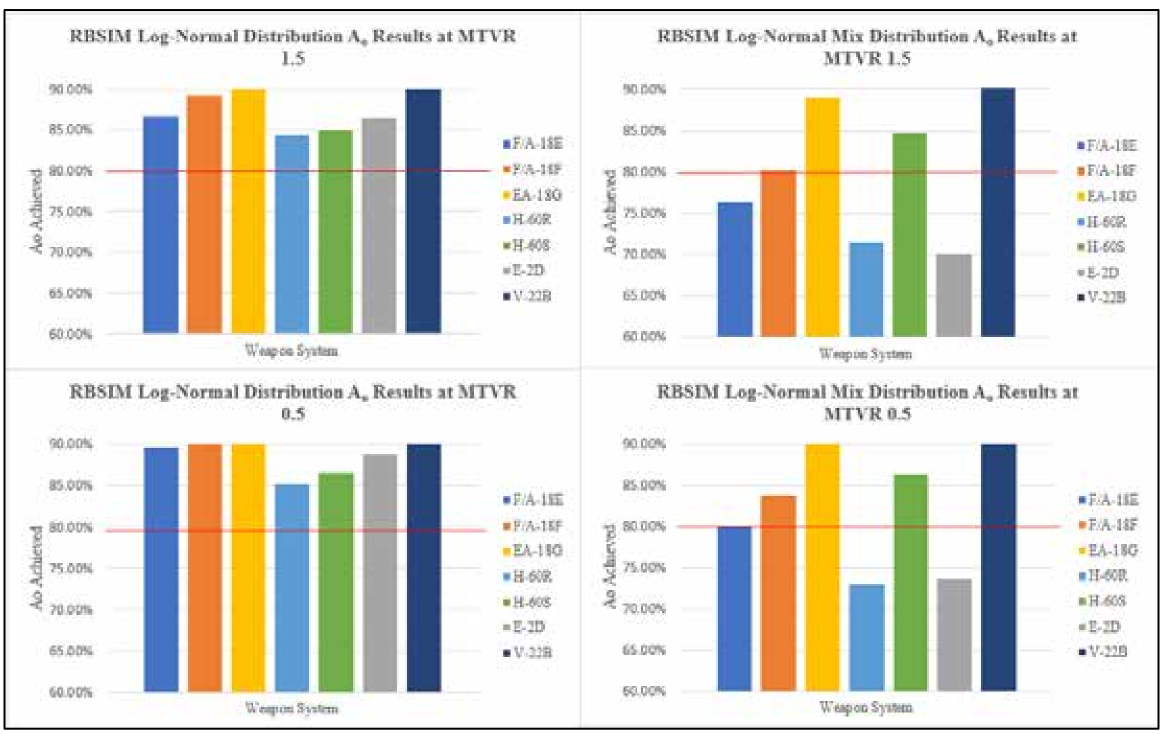

Figure 17. RBSIM Log-Normal Distribution A_o Results by WS.

Lastly, we group the distributions and evaluate the A_o based on the two MTVRs. Figures 18 and 19 shows the results. The graphs show that, consistently, the 1.5 MTVR achieves a lower A_o (averaging 29%) than that of the 0.5 MTVR (100%). Only two distributions achieve the target A_o with a 1.5 MTVR while all 0.5 MTVR achieve the target A_o. It appears the MTVR also makes difference on achieving the target A_o for the distribution.

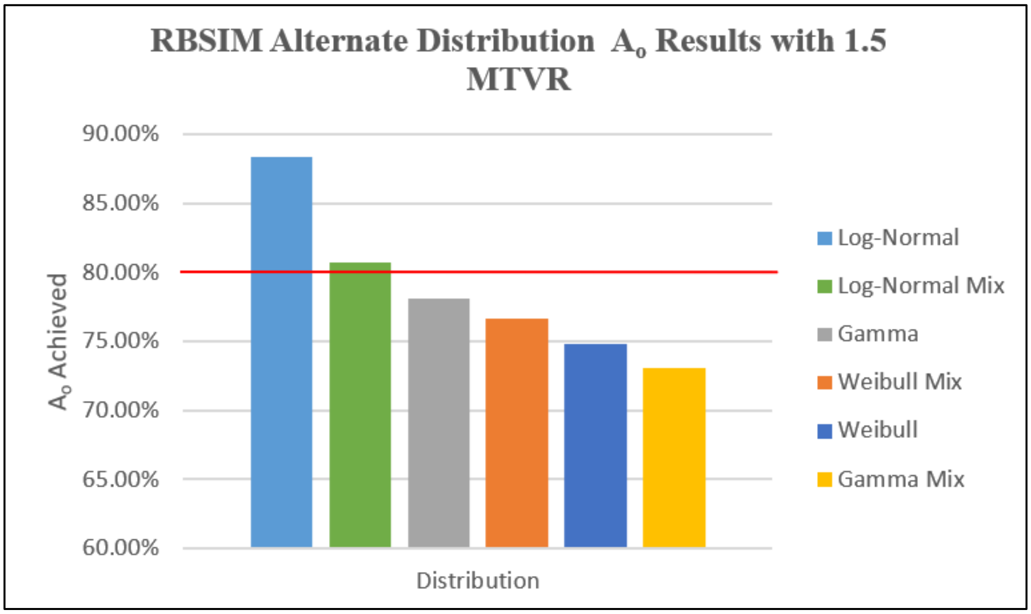

Figure 18. RBSIM Collective A_o Results on the Alternate Distributions at 1.5 MTVR.

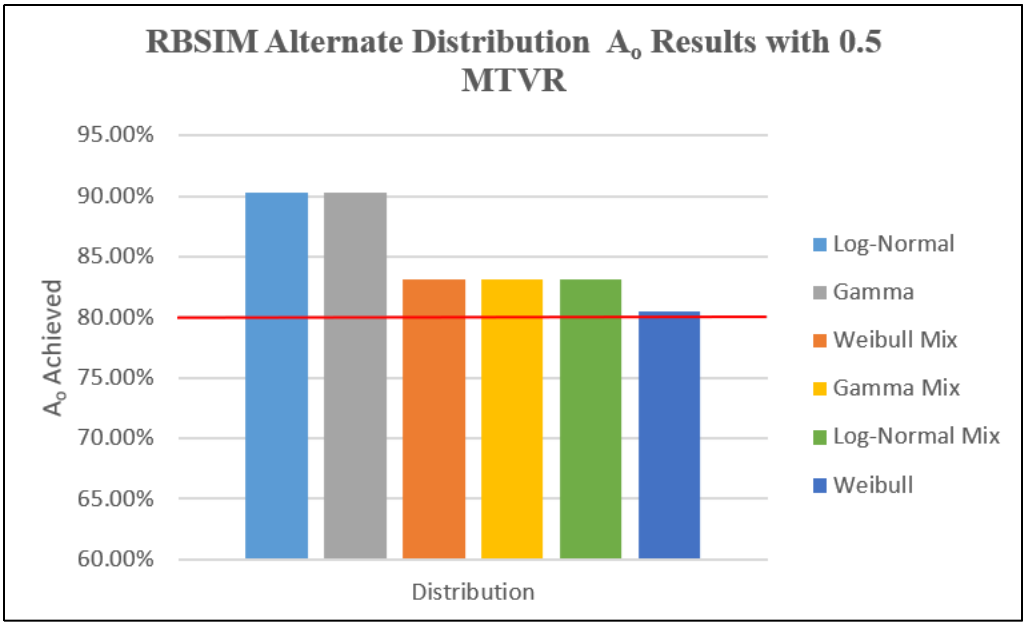

Figure 19. RBSIM Collective A_o Results on the Alternate Distributions at 0.5 MTVR

C. STATISTICAL ANALYSIS OF RBSIM A_O RESULTS

Finally, we perform a formal statistical analysis of RBSIM's results. For the analysis we use the R Statistical Software language. We capture a statistically test each RBSIM replication in order to understand the distribution of the A_o results. We use the RBSIM results from the F/A-18E, E-2D, and H-60R to analyze. These WSs have been chosen in order to sample one of each type of aircraft.

One goal is to test for normality in the RBSIM-generated A_o samples. A number of statistical tests, including the t-test we use later in this chapter, requires a normally distributed sample population. Verifying that RBSIM's A_o results are normally distributed allows us to continue with the statistical analysis. We accomplish this by taking the 30 replications of each WS estimated A_o and perform the Shapiro-Wilks test for normality (Taeger, 2015, p. 148). This test detects the departures from normality and rejects the null hypothesis of normality when the p-value is less than or equal to a specified value. We also perform two-sample t-tests comparing the estimated A_o means of the different alternate distributions and MTVRs (Taeger, 2015, p. 27). These t-tests consist of testing one WS using (a) two different distributions and different MTVRs; (b) the same distributions but different MTVRs; and, (c) two different distributions but the same MTVR. Here we are testing the null hypothesis that the true difference in the means are equal to zero.

We first test the F/A-18E WS. Each distribution with both MTVRs is run through the Shapiro-Wilks test for normality. The null hypothesis for this test is that RBSIM's A_o results are normally distributed while the alternate hypothesis is that the A_o results depart from normality. All of the alternate distributions including the mixed distributions have p-values higher than the 0.05 threshold and therefore there is not enough statistical evidence to reject the null hypothesis. RBSIM's A_o results for the F/A-18E all show a high confidence level of being (approximately) normally distributed.

The F/A-18E WS's estimated A_o normality test allows alternate distributions to be compared using the two sample t-test. The null hypothesis for this test is that the difference in the two sample A_o means equal zero, that is, that both means are identical. All p-values for the eighteen different combinations are less than 0.0001. Therefore, the difference in the 100% Weibull mean to all other selected distribution means is statistically significant. Next, we test the same distribution but different MTVRs. In all cases, we reject the null hypothesis that the difference in the estimated means of the compared distributions equals zero.

Finally, we compare different distributions but the same MTVR. We find that one of the eighteen combinations returns p-value greater than 0.05. The 0.5 MTVR gamma and log-normal two sample t-test returns a p-value of 0.7991 as shown in Table 3. Here, there is not enough statistical evidence to reject the null hypothesis of identical means for the 0.5 MTVR gamma and 0.5 MTVR log-normal distributions. For the F/A-18E, the difference in the mean for gamma and log-normal distributions with a 0.5 MTVR is not statistically significant.

Table 3. F/A-18E 100% Gamma Distribution Comparison to Other Distributions with Same MTVR.

F/A-18E 100% Gamma Distribution Comparison to The Other Distributions with Same MTVR Results		
Distributions	p-value	Null Hypothesis (H_o): $\mu_1 - \mu_2 = 0$
Gamma 1.5, Weibull Mixed 1.5	<0.0001	Reject H_o
Gamma 0.5, Weibull Mixed 0.5	<0.0001	Reject H_o
Gamma 1.5, Log-Normal 1.5	<0.0001	Reject H_o
Gamma 0.5, Log-Normal 0.5	0.7991	No evidence to reject H_o
Gamma 1.5, Log-Normal Mixed 1.5	<0.0001	Reject H_o
Gamma 0.5, Log-Normal Mixed 0.5	<0.0001	Reject H_o

The next WS we test is the E-2D WS's estimated A_o results for normality. All of the Weibull alternate distributions have a p-value below the 0.05 threshold and therefore there is not enough statistical evidence to reject the null hypothesis that Weibull's estimated A_o results are normally distributed. The E-2D's 100% alternate gamma distribution with the 1.5 MTVR has a low p-value for the Shapiro-Wilks test at 0.1450, which suffices to avoid rejecting the null hypothesis. The Shapiro-Wilks p-value for the 0.5 MTVR for gamma is slightly higher at 0.1505. In contrast with the 100% alternate distribution, the 1.5 MTVR mixed distribution has a much higher p-value for the normality test at 0.6990. The 0.5 MTVR mixed distribution also produces a higher normality test p-value of 0.6715.

The last distribution we test for normality for the E-2D WS is log-normal. The 1.5 MTVR 100% log-normal distribution satisfies the normality test. The resulting p-value is 0.5329. However, the log-normal distribution with a 0.5 MTVR rejects the null hypothesis for the normality test at 30 replications with a p-value of 0.0069 as Table 4 shows. Figure 20 shows the log-normal distribution of A_o. Figure 21 shows the QQ-plot of the same distribution. The gray area represents a normality reference line with the *y* axis representing the A_o results and the *x* axis representing the mean at zero with points of separation to both sides. As Figure 20 shows, both tails of the distribution depart from normal with majority of the replication points fall within the normality structure. The 100 independent runs of the 30 replications results are produced from RBSIM and the p-value for the 0.5 log-normal distribution is 0.9096. The 100 independent runs are required for the E2-D's log-normal distribution in order for the replications to become normally distributed. Figure 22 shows the results. Lastly, both log-normal mixed distributions pass the Shapiro-Wilks test with a p-value of 0.1522 for the 1.5 MTVR and a p-value of 0.6612 for the 0.5 MTVR. The majority of the E-2D WS A_o results do not reject the null hypothesis for the normality test. The single instance that initially did so is now (after 100 averages of the 30 replications are tested using the Shapiro-Wilk test) not rejecting the normality test.

Table 4. E-2D 100% Alternate Distribution Results for the Shapiro-Wilks Normality Test.

E-2D 100% Alternate Distribution Results for The Shapiro-Wilk Normality Test		
Distribution	p-value	Null Hypothesis (H_o): A_o results normally distributed
100% Weibull 1.5 MTVR	0.6690	Insufficient evidence to reject H_o
100% Weibull 0.5 MTVR	0.1450	Insufficient evidence to reject H_o
100% Gamma 1.5 MTVR	0.3547	Insufficient evidence to reject H_o
100% Gamma 0.5 MTVR	0.1435	Insufficient evidence to reject H_o
100% Log-Normal 1.5 MTVR	0.1505	Insufficient evidence to reject H_o
100% Log- Normal 0.5 MTVR	0.0069	Reject H_o

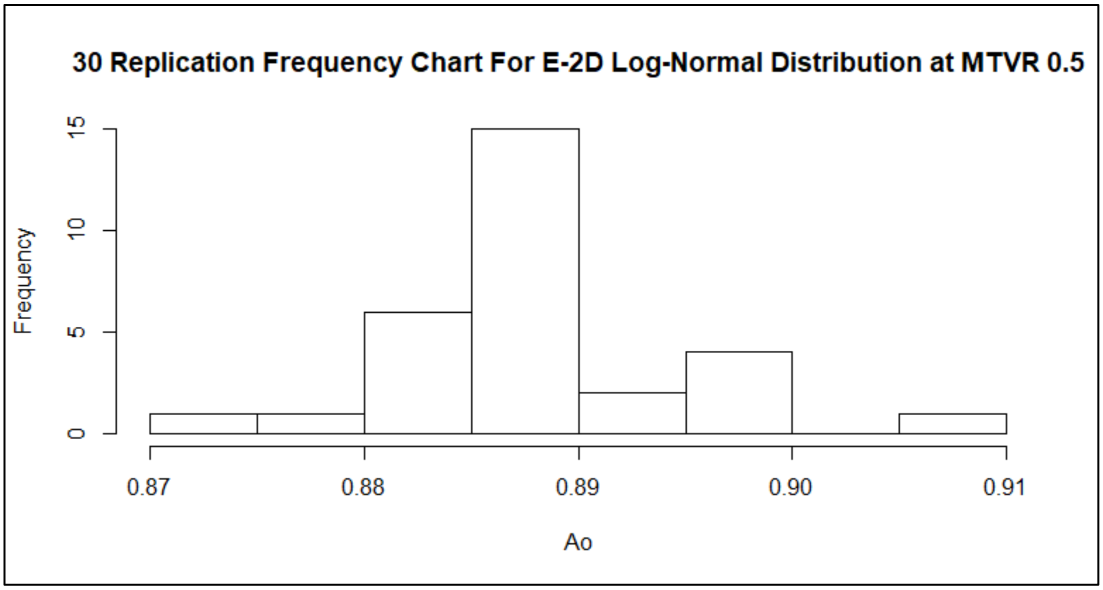

Figure 20. 30 Replications of E-2D Log-Normal Distribution with MTVR 0.5.

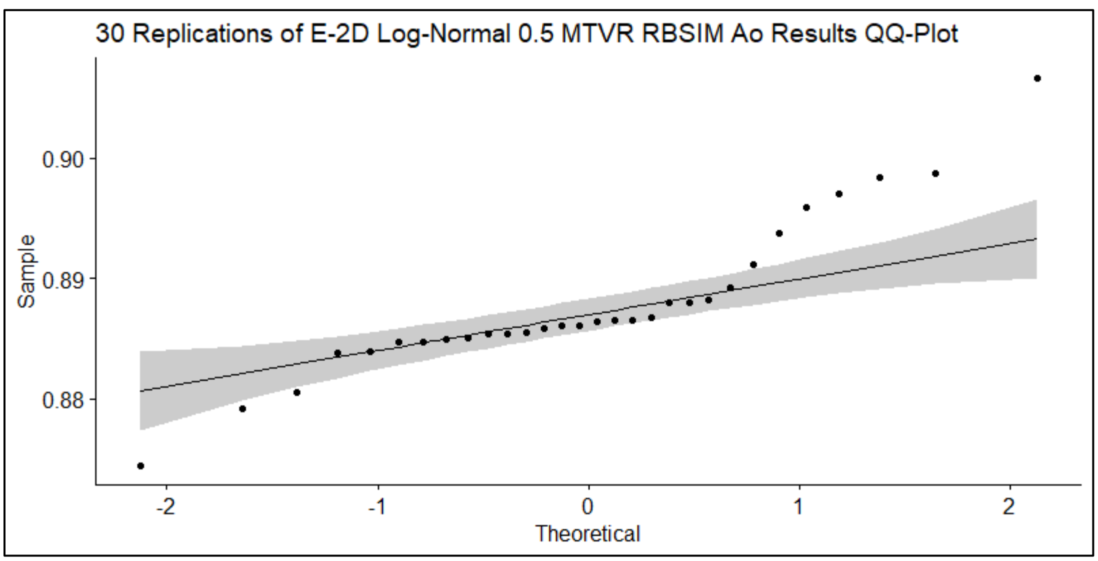

Figure 21. RBSIM E-2D Log-Normal Distribution A_o Results at 0.5 MTVR QQ-Plot.

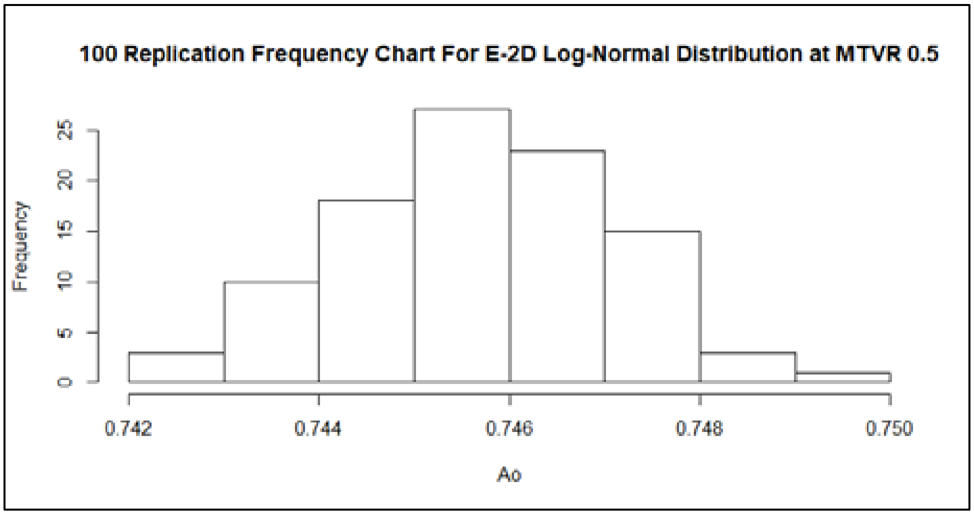

Figure 22. 100 Replications of Mean A_o for E-2D Log-Normal Distribution with MTVR 0.5.

The first t-test we perform is the comparison of the means of two different distributions with different MTVRs. Table 5 is the 100% Weibull distribution p-value results. The Weibull 0.5 and gamma 1.5 MTVRs t-test returns a p-value of 0.4861. There is not enough statistical evidence to reject the null hypothesis and therefore the difference in the means of these two distributions are not statistically significant. The remaining seventeen combinations have a p-value less than 0.05 and therefore the difference is statistically significant. The results for the same distribution but different MTVRs t-test produce the p-value less than 0.0001. All six combinations have a statistically significant difference in means. The final t-test comparing different distributions but the same MTVR returns one result with a p-value higher than 0.05. As Table 6 shows, the 0.5 MTVR gamma and log-normal distributions two sample t-test returns a p-value of 0.8971. There is not enough statistical evidence to reject the null hypothesis and the difference in the two means is not statistically significant.

Table 5. E-2D 100% Weibull Distribution Comparison Results to the Other Distributions with Different MTVRs.

E-2D 100% Weibull Distribution Comparison to The Other Distributions with Different MTVR Results		
Distribution	p-value	Null Hypothesis (H_o): $\mu_1 - \mu_2 = 0$
Weibull 1.5, Gamma 0.5	<0.0001	Reject H_o
Weibull 0.5, Gamma 1.5	0.4861	Insufficient evidence to reject H_o
Weibull 1.5, Log-Normal 0.5	<0.0001	Reject H_o
Weibull 0.5, Log-Normal 1.5	<0.0001	Reject H_o
Weibull 1.5, Gamma Mixed 0.5	<0.0001	Reject H_o
Weibull 0.5, Gamma Mixed 1.5	<0.0001	Reject H_o
Weibull 1.5, Log-Normal Mixed 0.5	<0.0001	Reject H_o
Weibull 0.5, Log-Normal Mixed 1.5	<0.0001	Reject H_o

Table 6. E-2D 100% Gamma Distribution Comparison to Different Distribution but Same MTVR Results.

E-2D 100% Gamma Distribution Comparison to The Other Distributions with Same MTVR Results		
Distribution	p-value	Null Hypothesis (H_o): $\mu_1 - \mu_2 = 0$
Gamma 1.5, Weibull Mixed 1.5	<0.0001	Reject H_o
Gamma 0.5, Weibull Mixed 0.5	<0.0001	Reject H_o
Gamma 1.5, Log-Normal 1.5	<0.0001	Reject H_o
Gamma 0.5, Log-Normal 0.5	0.8971	Insufficient evidence to reject H_o
Gamma 1.5, Log-Normal Mixed 1.5	<0.0001	Reject H_o
Gamma 0.5, Log-Normal Mixed 0.5	<0.0001	Reject H_o

We select the H-60R WS to test the rotary wing platforms. The H-60R WS's 100% alternate Weibull distribution with a 1.5 MTVR has one of the highest p-values among the WS samples at 0.8247 for the Shapiro-Wilks test. The p-value for Weibull's 0.5 MTVR decreases to 0.4950 but still rejects the null hypothesis for the normality test. When we apply it to the 1.5 MTVR mixed Weibull distribution we observe a p-value of 0.1093. Lowering the MTVR to 0.5 increases the p-value for the mixed distribution to 0.3908.

The 100% gamma distribution for the 1.5 MTVR also appears to be normally distributed. It has a p-value of 0.4310. The 0.5 MTVR 100% gamma distribution p-value result does reject the null hypothesis for normality at 30 replications. As Table 7 shows, the p-value is 0.0261. Figure 23 shows the histogram of the A_o distribution. However, Figure 24 shows that only the lower tail departs from normality while the rest of the results are in line with a normal distribution and with more replications it is highly likely that the distribution would be approximately normal. The 1.5 MTVR mixed distribution has the highest p-value for gamma's normality test at 0.7918 with the 0.5 MTVR gamma mixed distribution having the second highest p-value at 0.5496. The H-60R WS A_o results show that they appear to be approximately normally distributed and the single one that rejects the null hypothesis of a normal distribution only departs from normal at the lower end of the QQ-plot tail.

Table 7. H-60R 100% Alternate Distribution Results for the Shapiro-Wilks Normality Test.

H-60R 100% Alternate Distribution Results for The Shapiro-Wilk Normality Test		
Distribution	p-value	Null Hypothesis (H_o): A_o results normally distributed
100% Weibull 1.5 MTVR	0.8247	No evidence to reject H_o
100% Weibull 0.5 MTVR	0.4950	No evidence to reject H_o
100% Gamma 1.5 MTVR	0.4310	No evidence to reject H_o
100% Gamma 0.5 MTVR	0.0261	Reject H_o
100% Log-Normal 1.5 MTVR	0.6844	No evidence to reject H_o
100% Log- Normal 0.5 MTVR	0.6292	No evidence to reject H_o

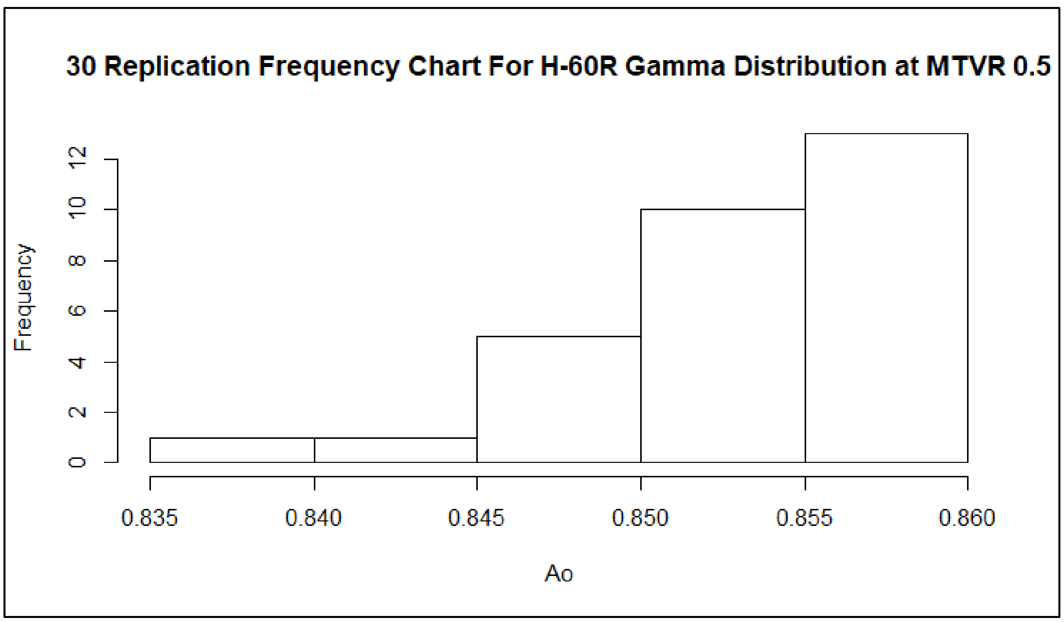

Figure 23. 30 Replications of H-60R Gamma Distribution with MTVR 1.5

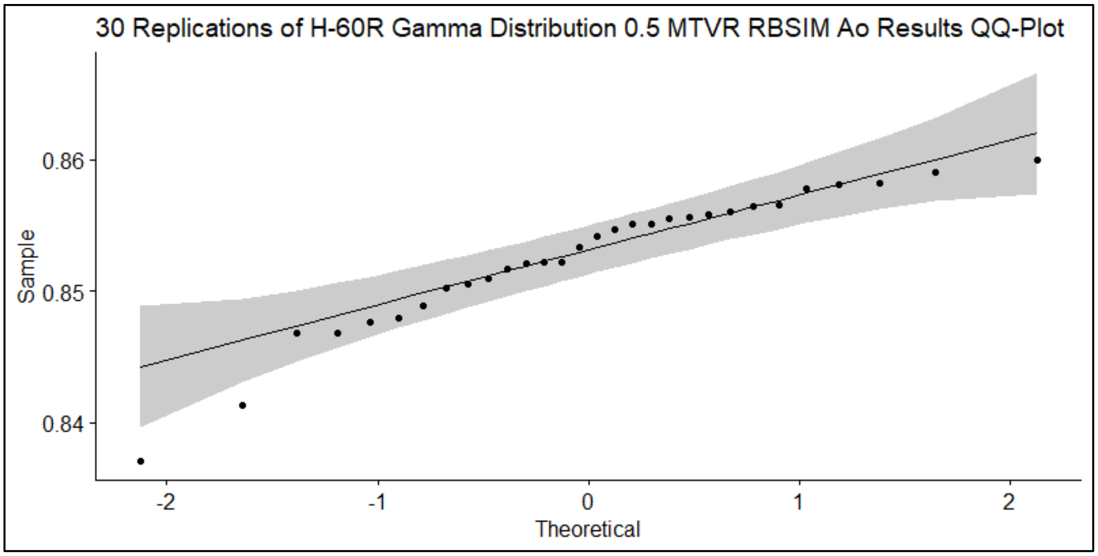

Figure 24. RBSIM H-60R Gamma Distribution A_o Results at 0.5 MTVR QQ-Plot.

As before, the first two sample t-test we perform for the H-60R's distributions is the comparison of the means of two different distributions with different MTVRs. Just like the E-2D's t-test for the same classification, the Weibull 0.5 and gamma 1.5 MTVRs two sample t-test returns a p-value 0.2123, higher than the 0.05 null hypothesis rejection threshold. The difference in the means of these two distributions is not statistically significant. All of the other combinations have a p-value less than 0.05 and therefore the means are not statistically the same. The t-test results for the same distribution but different MTVRs produce the p-value lesser than 0.0001 for all combinations therefore there is enough statistical evidence to reject the null hypothesis that the estimated A_o means are the same given the distribution selection criteria. The final t-test comparing different distributions but the same MTVR returns one result with a p-value higher than 0.05. The 0.5 MTVR gamma and log-normal distributions two sample t-test returns a p-value of 0.6496 as shown in Table 8. This is the same two sample t-test combination that returns a high p-value for the E-2D as well. There is not enough statistical evidence to reject the null hypothesis that the difference in these two sample means equal zero and therefore the difference is not statistically significant.

Table 8. H-60R 100% Gamma Distribution Comparison of Different Distribution with Same MTVR.

Results for H-60R 100% Gamma Distribution Comparison to The Other Distributions with Same MTVR		
Distribution	p-value	Null Hypothesis (H_o): $\mu_1 - \mu_2 = 0$
Gamma 1.5, Weibull Mixed 1.5	<0.0001	Reject H_o
Gamma 0.5, Weibull Mixed 0.5	<0.0001	Reject H_o
Gamma 1.5, Log-Normal 1.5	<0.0001	Reject H_o
Gamma 0.5, Log-Normal 0.5	0.6496	No evidence to reject H_o
Gamma 1.5, Log-Normal Mixed 1.5	<0.0001	Reject H_o
Gamma 0.5, Log-Normal Mixed 0.5	<0.0001	Reject H_o

V. CONCLUSION AND RECOMMENDATION

In this chapter we discuss our conclusions from the data analysis and offer some recommendations of further study of NAVARM.

A. CONCLUSIONS

RBSIM was initially developed by Wray to assess the validity of NAVARM's outputs. NAVARM calculates allowance levels to achieve the target A_o assuming an exponential demand distribution for all parts. These allowance levels are used to build an aircraft carrier's AVCAL. This thesis is testing the sensitivity of the distribution assumption made in NAVARM by introducing alternative distributions to the demand pattern. We do so on the basis that these distributions are known to possess statistical properties for modeling mean time between failures, for certain types of parts, and under certain assumptions. Based on the analysis performed in Chapter IV, we can conclude the following:

- Alternative distributions appear to have an impact on A_o.
- The MTVR also appears to impact A_o.
- The Weibull distribution produces the lowest A_o among the three alternative distributions. NAVARM's allowance levels are not high enough to meet target A_o under this demand distribution.
- NAVARM potentially does not allocate enough parts for the AVCAL where the MTBF distribution is non-exponential.
- The newest aircraft to the fleet, the EA-18G, achieves target A_o regardless of distribution.
- The log-normal distribution achieves the highest A_o among the three alternate distributions.

B. FUTURE WORK RECOMMENDATION

While this thesis concentrates on NAVARM's demand distribution assumption, further analysis would benefit the continued development of naval aviation RBS models. The following is a list of recommendations for future study in this area:

- Perform an in-depth analysis on RBSIM. Compare RBSIM A_o to actual A_o recorded from several sites to validate RBSIM output.
- Perform sensitivity analysis on any other NAVARM assumptions. For example, allow NAVARM to transfer parts between sites to cross-level inventory or allow for WS partial mission capable and record the impacts on A_o as well as allowance levels.
- Modify RBSIM to reflect real world supply and maintenance practices by introducing cannibalization practices, prioritized queues for resupply, and conditional failure rates by parts position and rerun this thesis's sensitivity analysis.
- Perform a statistical analysis of NAVSUP's candidate file with the goal of accurately capturing the real world MTBF in order to improve NAVARM's allowance calculation.
- Test the stationarity assumption of demands. In particular, measure the impact of a "surge" in demand on NAVARM's solution to whether or not the allowance levels can absorb the new demand and still achieve target A_o.
- In a similar vein, examine the impact of wartime operational-tempo on the AVCAL. The goal is to verify that NAVARM builds allowance levels using war-time flight hours that accurately capture wartime level of demand.

LIST OF REFERENCES

Buss, A. (2019). *Discrete event simulation (DES) modeling*, Working Paper, Naval Postgraduate School, Monterey, CA.

Chief of Naval Operations. (2003). *Operational availability of equipment and weapons systems. OPNAV instruction 3000.12A.* Washington, DC: Department of the Navy: Chief of Naval Operations.

———. (2011). *Readiness Based Sparing.* OPNAV Instruction 4442.5A. Washington, DC: Department of the Navy, Chief of Naval Operations.

Costantino, F., Di Gravio, G., & Tronci, M. (2013). Multi-echelon, multi-indenture spare parts inventory control subject to system availability and budget constraints. *Reliability Engineering & System Safety*, *119*, 95–101. https://doi.org/10.1016/j.ress.2013.05.006

Galliher, H., Morse, P., & Simond, M. (1959). Dynamics of two classes of continuous-review inventory systems. *Operations Research*, *7*(3), 362–384. https://doi.org/10.1287/opre.7.3.362

Li, Y., Feng, Y., Xue, X., & Lu, C. (2017). A united allocation method of spare parts and ground maintenance equipment for civil aircraft. *MATEC Web of Conferences*, *114*, 03006. https://doi.org/10.1051/matecconf/201711403006

Muckstadt, J. (1973). A model for a multi-item, multi-echelon, multi-indenture inventory system. *Management Science*, *20*(4-Part-I), 472–481. https://doi.org/10.1287/mnsc.20.4.472

Patriarca, R., Hu, T., Costantino, F., Di Gravio, G., & Tronci, M. (2019). A system-approach for recoverable spare parts management using the discrete Weibull distribution. *Sustainability (Basel, Switzerland)*, *11*(19), 5180–. https://doi.org/10.3390/su11195180

Sherbrooke, C. (1968). METRIC: A multi-echelon technique for recoverable item control. *Operations Research*, *16*, 122–141. https://doi.org/10.1287/opre.16.1.122

———. (1986). VARI-METRIC: Improved approximations for multi-indenture, multi-echelon availability models. *Operations Research*, *34*(2), 311–319. https://doi.org/10.1287/opre.34.2.311

———. (2004). *Optimal Inventory Modeling of Systems*. Springer.

Slay, F. M. (1984). VARI-METRIC: *An approach to modeling multi-echelon resupply when the demand process is poisson with a gamma prior.* Report AF301-3, Logistics Management Institute, Washington, D.C.

Taeger, D. (2015). *Statistical hypothesis testing with SAS and R.* John Wiley & Sons.

Wray, J. (2017). *A simulation of readiness-based sparing policies*. Master's thesis, Naval Postgraduate School, Monterey, CA.

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
   Ft. Belvoir, Virginia

2. Dudley Knox Library
   Naval Postgraduate School
   Monterey, California