



NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

**MEASURING THE IMPACT OF BUSINESS RULES ON
INVENTORY BALANCING**

by

Andrew Oswald

September 2013

Thesis Co-Advisors:

Michael P. Atkinson

Geraldo Ferrer

Second Reader:

Walter DeGrange

Approved for public release; distribution is unlimited

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, 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 Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE (DD-MM-YYYY) 23-9-2013			2. REPORT TYPE Master's Thesis		3. DATES COVERED (From — To) 2011-07-05—2013-09-27	
4. TITLE AND SUBTITLE MEASURING THE IMPACT OF BUSINESS RULES ON INVENTORY BALANCING					5a. CONTRACT NUMBER	
					5b. GRANT NUMBER	
					5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Andrew Oswald					5d. PROJECT NUMBER	
					5e. TASK NUMBER	
					5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943					8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Department of the Navy					10. SPONSOR/MONITOR'S ACRONYM(S)	
					11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited						
13. 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 XXX.						
14. ABSTRACT Naval Supply Systems Command recently employed the Navy Enterprise Resource Planning Single Supply Solution to improve efficiency through the sharing of data across the organization. For the first time, replenishment decisions were made using shared enterprise data. Since all items from all sites are in one central database, Weapon Systems Support has total visibility of assets across all available supply sources. Inventory balancing is a promising functionality for enhancing the performance of inventory systems. With a balancing policy in place, stock can be moved from a location that has excess inventory to another location experiencing a shortage. The purpose of this movement of materiel is to reduce inventory costs and increase the percentage of demand satisfied by on-hand stock. This paper compares the relative effectiveness of different balancing business rules. A simulation model comprising a two-echelon supply network with three warehouse locations is used to evaluate the various business rules for several items with varying unit prices and demand frequencies. Although no single balancing policy is optimal in all situations, simple modifications to the proposed business rules will increase the benefits of balancing while minimizing any negative effects.						
15. SUBJECT TERMS Inventory balancing; simulation; supply chain management; inventory control; Lateral transshipment; Spare parts; Inventory pooling; Navy Enterprise Resource Planning						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON	
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (include area code)	
Unclassified	Unclassified	Unclassified	UU	113		

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited

MEASURING THE IMPACT OF BUSINESS RULES ON INVENTORY BALANCING

Andrew Oswald
Lieutenant Commander, Supply Corps, United States Navy
B.S., United States Naval Academy, 2001

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN OPERATIONS RESEARCH

from the

**NAVAL POSTGRADUATE SCHOOL
September 2013**

Author: Andrew Oswald

Approved by: Michael P. Atkinson
Thesis Co-Advisor

Geraldo Ferrer
Thesis Co-Advisor

Walter DeGrange
Second Reader

Robert F. Dell
Chair, Department of Operations Research

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

Naval Supply Systems Command recently employed the Navy Enterprise Resource Planning Single Supply Solution to improve efficiency through the sharing of data across the organization. For the first time, replenishment decisions were made using shared enterprise data. Since all items from all sites are in one central database, Weapon Systems Support has total visibility of assets across all available supply sources. Inventory balancing is a promising functionality for enhancing the performance of inventory systems. With a balancing policy in place, stock can be moved from a location that has excess inventory to another location experiencing a shortage. The purpose of this movement of materiel is to reduce inventory costs and increase the percentage of demand satisfied by on-hand stock. This paper compares the relative effectiveness of different balancing business rules. A simulation model comprising a two-echelon supply network with three warehouse locations is used to evaluate the various business rules for several items with varying unit prices and demand frequencies. Although no single balancing policy is optimal in all situations, simple modifications to the proposed business rules will increase the benefits of balancing while minimizing any negative effects.

THIS PAGE INTENTIONALLY LEFT BLANK

Table of Contents

1	Introduction	1
1.1	Background	1
1.2	Literature Review	9
1.3	Purpose of the Analysis.	17
1.4	Scope, Limitations, and Assumptions	17
1.5	Research Questions	18
2	Data and Methodology	19
2.1	Data	20
2.2	Model Formulation	27
2.3	Simulation Structure and Rationale	36
3	Analysis	49
3.1	Initial Results.	50
3.2	Sensitivity Analysis	62
3.3	Summary	76
4	Conclusion and Recommendations	79
4.1	Discussion	79
4.2	Recommendations to Improve the Balancing Policy.	84
4.3	Future Research.	85
	References	87
	Initial Distribution List	89

THIS PAGE INTENTIONALLY LEFT BLANK

List of Figures

Figure 2.1	Bar Chart of Excess and Deficit Inventory Status	21
Figure 2.2	Pie Chart of Excess Inventory by Line Items and Dollar Value	22
Figure 2.3	Venn Diagram of Commonality of Demand	23
Figure 2.4	Line Plot of Distribution by Value of National (or NATO) Item Identification Numbers (NIINs)	24
Figure 2.5	Bar Charts Contrasting Demand Data for Two Items at Two Locations	27
Figure 2.6	Scatter Plot of Weight and Shipping Cost	28
Figure 2.7	Scatterplot of Shipping Cost Versus Weight with Regression Line . . .	31
Figure 2.8	Histogram of Demands Per Day	39
Figure 2.9	Flowchart of Outstanding Order Arrival Event	41
Figure 2.10	Flowchart of Demand Event	41
Figure 2.11	Flowchart of the Reset Event	42
Figure 2.12	Flowchart of Placing Order Event	43
Figure 2.13	Flowchart of Stage 1 of Redistribution Event	44
Figure 2.14	Flowchart of Stage 2 of Redistribution Event	45
Figure 2.15	Flowchart of Stage 3 of Redistribution Event	46
Figure 2.16	Flowchart of Stage 4 of Redistribution	47
Figure 2.17	Flowchart of Single Simulation Day for a Single Location	48
Figure 3.1	Bar Chart of Total Cost	53

Figure 3.2	Diagram of Excess Fulfilling Demand at a Location with MSL Equal to Zero	57
Figure 3.3	Bar Chart of Monthly Demands at Oceana for the VALVE, REGULATING item	59
Figure 3.4	Bar Chart of Fill Rate and Inventory Costs Before and After Balancing for Various Amounts of Initial Excess	64
Figure 3.5	Bar Chart of Fill Rate and Inventory Costs Before and After Balancing for Various Amounts of Initial Excess with an Increase in Allowance (Patuxent River)	65
Figure 3.6	Bar Chart of Fill Rate and Inventory Costs Before and After Balancing for Various Amounts of Initial Excess with an Increase in Allowance (Oceana)	66
Figure 3.7	Venn Diagram of Commonality of Demand	67
Figure 3.8	Diagram of Excess Fulfilling Demand at a Location with Comparably High Allowance	71

List of Tables

Table 1.1	Six Categories of Navy Supply Management	7
Table 1.2	Balancing Regions and Price Thresholds	7
Table 1.3	Storage Costs for materiel	13
Table 2.1	Excess Inventory Values for Lemoore, Oceana, and Patuxent River . . .	21
Table 2.2	Common Items Stratified by Unit Price Thresholds	23
Table 2.3	Patuxent River Listing of NIINs by Descending Dollar Demand	25
Table 2.4	Summary Data for Five Largest Items by Dollar Volume	25
Table 2.5	Observed Paired Data for Weight, Cost, and Predicted Cost	30
Table 2.6	Inventory Cost Model Variables, Values, and Source	35
Table 2.7	Simulation Events	38
Table 2.8	Poisson Distribution Description	40
Table 3.1	Cost and service impact of inventory balancing.	51
Table 3.2	Cost and Benefits of Inventory Balancing	54
Table 3.3	Fill Rates for NIIN 011669450	57
Table 3.4	Fill Rates for NIIN 012030382	60
Table 3.5	Fill Rates for NIIN 012030382 (Increased Level Setting)	61
Table 3.6	Percent of Total Cost Due to Stock Transfers	62
Table 3.7	Common Items Stratified by Unit Price Thresholds	67
Table 3.8	Summary Data for Five Largest Items by Dollar Volume	68

Table 3.9	Impact of Implementing Balancing	69
Table 3.10	Allowance and Two-Year Demand	70
Table 3.11	Fill Rate for NIIN 012030382 with Alternative ROP Policy	72
Table 3.12	Fill Rates for NIIN 011223338	73
Table 3.13	Allowance Changes	74
Table 3.14	Impact of Implementing Balancing	75
Table 3.15	Daily On-hand Inventory Balance and Procurements for NIIN 011223338	76
Table 3.16	Percent of Total Cost Due to Stock Transfers	76

List of Acronyms and Abbreviations

BP 28	Budget Project 28
CV	coefficient of variance
DLA	Defense Logistic Agency
DOD	Department of Defense
DON	Department of the Navy
EOQ	economic order quantity
ERP	Navy Enterprise Resource Planning
FIAR	Financial Improvement and Audit Readiness
FLC	Fleet Logistic Centers
GAO	Government Accountability Office
GLS	Global Logistics Support
GSA	General Service Administration
MSL	maximum stock level
NAS	naval air station
NAVSUP	Naval Supply Systems Command
NWCF	Navy Working Capital Fund
NIIN	National (or NATO) Item Identification Number
NLS	no lateral shipments
NSN	National Stock Numbers
OSPA	Operational Supply Planning Analyst
PO	purchase order
PR	purchase request
R-Supply	Relational Supply
RHICS	Regional Hazardous Inventory Control System

ROP	reorder point
RRAM	Real-time Reutilization Asset Management
SNI	single national inventory
SPR	Site Planning Representative
TRF	Trident Refit Facilities
U2	Uniform Automated Data Processing System for Stock Points - 2
U.S.	United States
WSS	Weapon Systems Support

Executive Summary

The Navy Enterprise Resource Planning (ERP) Single Supply Solution drastically increases the potential for improved inventory management efficiency through the sharing of data across an organization. Now including a centrally managed inventory system, a particular feature of the ERP, inventory balancing, is well-suited to addressing the \$794 million dollars of excess materiel held by the United States Navy. The goal of leveraging inventory balancing is to reduce this excess materiel as well as inventory management costs. However, inventory balancing is not being used. To maximize this functionality, Weapon Systems Support (WSS) has proposed a set of business rules. The Navy is concerned that the implementation of inventory balancing may lead to excessive churn and decreases in materiel availability. This is because a qualitative analysis has not been conducted in relation to the proposed business rules' ability to increase the percentage of demand satisfied by on-hand stock and lower inventory costs. In this study, we simulate the proposed business rules and measure their impact on fill rates and inventory management costs relative to not using balancing and variations of the proposed business rules.

Inventory balancing is beneficial when it moves materiel from a location that is unlikely to experience demand to a location that is more likely to experience demand. Business rules govern shipments between locations in an attempt to ensure that materiel move to where they are needed the most. These rules use the on-hand inventory levels relative to other factors to determine if the shipping of materiel to another location will improve fill rates and decrease inventory costs. Over a two-year period, the study compares the expected total costs and fill rates with and without balancing. Based on the outcome of the comparison, modifications to the business rules are proposed, tested, and compared. From this, we recommend modifications to the business rules to maximize the benefits and minimize the costs of inventory balancing.

In order to realize cost saving by efficiently using on-hand inventory, excesses and shortages must exist and there must be common demand among locations. The Navy owns hundreds of millions of dollars in excess inventory while simultaneously facing stock shortages. Excess inventory is defined as the number of units in stock above the maximum stock level (MSL), or order-up-to level, while a shortage is defined as the amount of inventory below the MSL. Given the inventory data from WSS for the naval air stations (NASs) at Lemoore, Oceana, and Patuxent River, we identify potential for the implementation of an inventory balancing policy. The data from these locations show that they experience both excesses and shortages, and share demand of common items. We use the top five from these commonly demanded items, according to dollar

volume, to test the impact of balancing in our simulation. We also find the common demand and excess indicative of potential for inventory balancing at Lemoore, Patuxent River, and Kings Bay, a submarine base. Again, we use the top five commonly demanded items among these sites to compare business rules.

We only consider relevant inventory management costs, which are those that change in value depending on the set of business rules employed. Any cost that remains the same regardless of the balancing policy is not relevant because it will not impact the decision to employ a particular set of business rules. The relevant costs include procurement, stock-out, inventory, backorder, and transshipment costs. To calculate the associated inventory management costs, we develop a mathematical model. The mathematical model calculates costs for a given amount of items. For example, replenishment costs depend on the number of procurements. The more the procurements, the higher the replenishment costs will be. Likewise, the more stock-outs that occur, the higher the stock-out costs will be, and so on.

To compute the associated costs, the model requires the number of procurements, stock-out occurrences, backorders, and transshipments occurrences as well as how much inventory is in stock each day. However, historical balancing data are not available given that a balancing policy has not yet been implemented. Analytic solutions rely on either constant or convenient demand and lead time distributions. These assumptions simplify the inventory analysis to make addressing this issue more practical. However, these simplifications misrepresent the NASs demand data. Instead, a computer simulation generates the daily procurement, stock-out, inventory, backorder, and transshipment costs using the mathematical model to calculate these costs per period. Analysis of these costs provides a mechanism through which to compare the cost efficiencies gained or lost using a given set of inventory balancing business rules.

The simulation provides several two-year estimates of possible “results” within minutes, which would not be feasible through experimentation. Any experiments using the physical system would require years for the data necessary to measure the possible impact of competing balancing procedures to be collected. In addition, a simulation does not interfere with the real system. Whereas there are inherent risks and costs associated with a real-world experiment, excessive costs within a simulation have no real-world consequences. Furthermore, this study adheres to the specific request by WSS to provide an analysis prior to physically implementing a balancing policy.

From the simulation results, we find that inventory balancing neither increases fill rates nor

decreases inventory management costs in all situations. In fact, fill rates actually decrease more often than they improve as a result of balancing in the examined instances. In addition, in the majority of cases, it decreases by less than one percent. To assess the balancing rules, we vary the location of 30 days of excess materiel. The 30-day quantity is based on the location with the highest demand. We first consider the impact of balancing when the initial excess resides at the high-demand location, then the moderate-demand location, and finally the low-demand location. Balancing performs poorly when all the excess resides in the high-demand location relative to when the excess resides in the locations with lower demand. The two-year savings from balancing for the nine unique items examined in this study when the excess resides at the high-demand location are \$5,322. When the excess resides in the low-demand location, the savings from balancing for the unique examined items are \$15 million.

In reality, the excess is probably distributed somewhere between these two extremes. However, we know that the amount of excess at Lemoore, Oceana, and Patuxent River is valued at \$43.3 million. It is reasonable to expect that some of this excess could offset replenishment costs through inventory balancing. These values provide a reasonable order of magnitude of the potential savings from inventory balancing. Although we only consider nine unique items, we choose those that are the most commonly demanded according to dollar volume. The inclusion of more items may increase potential savings, but their contributions will become marginally lower as their respective dollar volumes decrease.

Although no single balancing policy clearly improves both fill rates and inventory management costs in all situations, our analysis uncovered key insights that influence transshipment decisions in relation to whether to balance. The following circumstances influence when balancing is preferable over not balancing.

1. Balancing improves fill rates and costs when the amount of excess is greater than the future demand.
 - The impact of balancing is the greatest when the excess resides in the location without demand. It reduces holding and procurement costs.
 - When the amount of excess is less than the mean 90-day demand, the impact from balancing is minimal and, in some cases, detrimental.
2. When a location's on-hand inventory is zero, the shipment of excess to this location does not improve fill rates but can reduce inventory management costs.
 - Raising the MSL at a location by increasing its allowance from zero to two can

reverse this trend.

- Increasing the allowance from zero to two can also improve inventory management costs by reducing backorder costs.
3. Erratic demand can result in a low MSL when future demand is high. Inventory balancing exacerbates this issue by shipping excess away from where the demand is more likely to occur, thus decreasing fill rates, which increases backorder costs.
 - Increasing the frequency of level setting reverses this trend.
 - Performing inventory balancing only after level setting but before procurement also mitigates this issue.
 4. When a location has a relatively high allowance compared to demand (e.g., an allowance that is more than twice the mean two-year demand), using the MSL as an eligibility criterion to receive excess makes balancing less preferable.
 - By using the reorder point (ROP) instead of the MSL eligibility criterion to receive excess reduces and in some cases eliminates the negative impact of balancing that results when one location has a high allowance relative to demand.
 - Resetting the allowance to a more reasonable level greatly reduces actual costs, regardless of whether a balancing policy is employed. However, balancing represents an improvement over not balancing.
 5. Transshipment costs contribute less than three percent to inventory management costs in all situations. In none of the cases did transshipment costs determine whether balancing was beneficial.

From these insights, we propose the following modifications to the business rules to mitigate any negative outcomes from balancing while emphasizing its benefits.

1. Use dollar volume to determine balancing eligibility. Limit balancing to the 10 percent of items that account for the majority of inventory management costs. This will reduce the number of shipments between locations by only transshipping the items that will most likely result in reduced inventory costs.
2. Instead of balancing before generating purchase requests twice a week, only run inventory balancing after quarterly level setting, which occurs every 90 days. Although this will reduce the potential savings from inventory balancing, it will also reduce its potential risks. This recommendation will decrease the number of transfers between locations, ensure that materiel is packaged together, and decrease some of the observed negative impacts of inventory balancing. We find that increasing the frequency of level setting reverses the

negative impact on costs and fill rates under a balancing policy. Increasing the frequency of level setting may not be feasible, but limiting balancing to occur only after quarterly level setting results in similar benefits.

3. Remove balancing limitations between the geographic regions within the continental United States. The shipping times and prices between Oceana and Patuxent River are the same as those between Oceana and Lemoore. From the inventory cost and materiel availability perspectives, these geographic limitations are not relevant.
4. Do not transfer materiel to locations with on-hand inventory levels below their MSL. Instead, transfer inventory to locations with on-hand inventory levels below their ROP. This limits the potential negative impact of balancing, as seen in our analysis section, while maintaining most of its benefits.
5. Do not transfer excess materiel from locations if the amount of excess is lower than the 90-day demand forecast. When the excess resides in a location in which demand is relatively high compared to the amount of excess, the impact of balancing is minimal and, in some cases, detrimental.

THIS PAGE INTENTIONALLY LEFT BLANK

Acknowledgements

I would like to thank the exemplary operation research analysts and operational planners from Naval Supply Systems Command Weapon Systems Support and Global Logistics Support who assisted with the research in this study. **John Cardillo, Eric Liskow, LT Josh Lasater, Donald McNeeley, Brian Palko, LCDR Ernesto Ureta, and Charles Yahn**, your ability to provide data and answer my questions, in addition to the important work you already do on a daily basis, is astounding. Again, thank you for gathering all the data I needed for this thesis. I could not have done it without your help, and I want to commend you for your efforts.

My special thanks to **Jeff Powell** and **CDR Xavier Lugo**. Our early correspondence set the course for this thesis. Without their knowledge and contacts, this thesis would not exist. In addition, I would also like to thank **Charlie Albright** and **Cathy Wiechelt** for providing a perspective on Budget Project 28 materiel unobtainable from the data alone. **LCDR Bruce Reilly** and **LCDR Andr s Diaz**, thank you for your assistance during my experience tour.

Thank you to **CAPT David Meyers** for giving me the opportunity to pursue this thesis. I would also like to thank my advisors and second reader for their extreme patience. **Dr. Michael Atkinson, Dr. Geraldo Ferrer, and CDR Walter DeGrange**, my sincere thanks to all of you for making such a great effort to strengthen this thesis through your inputs, recommendations, and insights. Not only did you turn around edits in record time, but you provided impressively thorough comments and direction.

Thanks to my fellow OR students for their support, feedback, and friendship. I would like to thank my friends, especially “**Duncan, Ben, Spencer, Kevin, Liz, and Ashley**” for teaching me so much through our group projects, classes, and volleyball.

I would like to thank my family for supporting me and giving me a chance to thrive, especially my wife, **Allison**, for her feedback on various aspects of this thesis. There is no way I could complete this without her help and understanding.

Last, but certainly not least, I would like to thank the operations research faculty at the Naval Postgraduate School for their strength and dedication to the operations research profession.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 1:

Introduction

Since March 2010, the Naval Supply Systems Command (NAVSUP) has employed its Navy Enterprise Resource Planning (ERP) Single Supply Solution to improve efficiency by sharing data across the organization. This study investigates one particular feature of ERP, the inventory balancing function, which is not implemented, but has the potential to improve efficiency and save costs. As of November 2012, the Navy held over \$794 million in excess inventory. This figure could potentially be reduced by implementing inventory balancing, realizing significant cost savings. The private sector uses inventory balancing with great success to minimize costs and increase the percentage of orders satisfied from stock on hand. However, the Navy does not feel it fully understands the effects of the business rules that define inventory balancing, and so has indefinitely postponed implementing this functionality. This hesitation is the result of concerns about increases in shipping costs and the number of items that need to be shipped during balancing. This study provides an initial analysis of the impact of competing inventory balancing business rules. Decision makers can consider this information when assessing whether to implement the ERP balancing functionality. If a set of balancing business rules results in a lower total cost and fulfills the fleet's demand more often than without balancing, then implementing this function will benefit the Navy.

The existing Navy ERP system implementation has meant that inventory balancing has only recently become technologically possible. An investigation of the transition from the legacy systems to ERP will demonstrate why now is the time for decision makers to assess inventory balancing business practices. The inventory replenishment process is discussed in detail to help explain the analysis methodology. Existing literature on inventory balancing is also examined to define the terminology and to provide a common theoretical framework for this study. Answering the proposed research questions, given the study's limitations and scope, will provide decision makers with the necessary information to choose the inventory balancing rules that best meet their requirements.

1.1 Background

For years, the Navy has tried to unify its supply enterprise under a single inventory system, but has struggled to overcome the challenges inherent in incompatible systems and complex

business practices. These limitations resulted in the Government Accountability Office (GAO) designating the Department of Defense Department of Defense (DOD) financial management as high risk in 1995 (United States Government Accounting Office 1995). The GAO found that the DOD management systems were not able to provide accurate and timely financial management visibility and effective accountability. The divergent data management systems implemented at the time were ill-suited to providing an accurate audit trail, as each system recorded and stored data differently. The time needed to consolidate this data exceeded the reporting requirements for government oversight.

To address these concerns, the DOD initiated the Financial Improvement and Audit Readiness (FIAR) plan to define a strategy and the processes required to provide accurate and timely financial reports. A key aspect of this strategy is an enterprise resource planning tool, which is a software package that unifies the financial and resource management functions of the organization. Under the FIAR plan, each service is responsible for creating its own resource planning tool to meet the reporting requirements.

The Navy's first attempt to develop a resource planning tool, in 1998, was a billion-dollar failure (United States Government Accountability Office 2005). This effort failed because the Navy initiated four independent pilot programs, each addressing different areas of the Navy Enterprise. The intention was to integrate these pilot programs into a unified program. However, each program made design choices based on their area of the Navy Enterprise, so creating quite different processes for similar functions. In the end, the Navy found it impossible to combine the four pilot programs into one uniform solution.

The Navy built the next iteration of its ERP system from the ground up, as the previous programs were unsalvageable. A major difference of this effort was the appointment of a central programming office to oversee the development of the program and to focus on addressing the GAO's concerns regarding financial accountability (United States Government Accountability Office 2005). As a result, the Navy ERP system intends to standardize, not just financial management, but acquisition, maintenance, procurement, plant and wholesale supply, and workforce management within the Navy as well. This study focuses on plant and wholesale supply.

Prior to the ERP system, the Navy relied on small, individualized systems for wholesale and retail supply. Although the Navy was able to operate using these legacy systems, the process was complicated and inefficient, creating many layers with duplicate records. These overlapping systems impeded the Navy's ability to provide accurate and timely data. The ERP system re-

placed the Real-time Reutilization Asset Management (RRAM) system, the Regional Hazardous Inventory Control System (RHICS), the Uniform Automated Data Processing System for Stock Points - 2 (U2), and the Relational Supply (R-Supply) system at each naval air station (NAS). The four legacy systems were completely retired from the NAS as of September 2012.

One of the key elements of the Navy ERP implementation is the creation of a Navy-wide single national inventory (SNI). Created to meet the National Inventory Management Strategy (NIMS) initiative, the SNI unifies the once disparate inventories within the DOD into a single inventory managed by an organization-wide system, even though it is not housed in a single location. Prior attempts to establish an SNI failed, as the decentralized legacy systems could not overcome the difficulties of asset visibility over multiple databases. By replacing several disjointed legacy IT systems, the Navy ERP system has a single database across the entire enterprise, which provides the functionality necessary to manage an SNI.

Another factor that contributed to the failure of the SNI under the legacy policy was the supply planning function, the process of deciding to buy or repair inventory, which occurred at individual sites rather than on a Navy-wide basis. The Navy ERP system standardizes and centralizes the inventory control processes. By centralizing the data and decision-making process, the Navy can capitalize on the pooled expertise of the NAVSUP Weapon Systems Support (WSS) instead of relying on individual site representatives. By managing all the sites centrally, decision makers have an enterprise-wide view that enables them to manage inventory more effectively. However, an unintended side effect of centralizing supply planning is the loss of the local perspective, such as shifts in maintenance plans for a given platform (e.g. helicopter, ship) that may require more parts than previously needed.

To ensure that both local and Navy-wide perspectives are incorporated, two new roles have been included in the Navy ERP to ease the transition. The Operational Supply Planning Analyst (OSPA) in WSS Philadelphia will provide overall inventory management for Navy-managed inventory in the ERP system. A symbiotic relationship between the OSPA and the Site Planning Representative (SPR), located at the site where inventory was once managed, provides balance in the process. The SPR provides local insight into demand plans, forecasts, allowances, inventory levels, and replenishment plans to the OSPA. Together, the two roles retain the site-level perspective for demand planning while still benefitting from consolidated data.

With accurate and timely reporting across all sites, replenishment decisions can consider Navy-wide shortages and excesses. The catalyst to shortages and excesses at sites is demand variation.

Two business practices address variation in demand. The first, demand planning, uses forecasted demand to set the maximum and reorder inventory levels (i.e., the on-hand inventory). The second, supply planning, redistributes and replenishes inventory to support the new levels set by demand planning. Since the item data from all sites are in one central database, the WSS can see all available supply sources, and any excess inventory at one site can be redistributed as required. Under the legacy systems and policy, this was impossible.

The Navy holds inventory to maintain and repair military assets to meet operational requirements in the face of fluctuating demand, to avoid situations in which an item cannot be fulfilled locally, and to shift repairs and maintenance schedules when costs are relatively low. The Navy ERP system enables inventory to be redistributed across sites to help maintain optimum inventory levels. Holding too much inventory is undesirable because the excess ties up funds that could be used elsewhere within the Navy. Inventory also incurs a cost, as the stored inventory needs to be cared for and is subject to obsolescence. However, not carrying any inventory, or holding too little inventory, requires more frequent replenishments, as well as increased shipping, handling, and ordering costs. Holding inventory is necessary to meet government fill rate mandates (i.e., the percentage of demand satisfied from stock on hand), and low levels of inventory may result in a shortage of assets necessary for the Navy's operation.

Determining the "right" inventory level requires balancing the advantages of carrying more stock, including higher fill rates and bulk discount pricing at the time of purchase, against the disadvantages, such as tying up funds and running the risk of inventory obsolescence at each plant. The plants that store inventory include the NAS, the Trident Refit Facilities (TRF), and the Fleet Logistic Centers (FLC). Using historical data, the Navy ERP system forecasts demand for each month and each site by individual National Stock Numbers (NSN). This demand planning determines the level of stock that should be maintained at each site.

Demand planning establishes demand-based levels for Navy-managed inventories. The aforementioned forecasts are used to set order-up-to inventory levels. The order-up-to inventory level, or maximum stock level (MSL), is the target inventory level for an item. When the inventory level reaches a set reorder point (ROP), an order is placed to return the on-hand stock quantity to this target value. The MSL is large enough to ensure that stock on hand can cover the lead time between the scheduled reorder and delivery of the stock (Silver et al. 1998). As mentioned, the forecasts extrapolate historical demand data, but also incorporate human judgment of future requirements (Silver et al. 1998). Although the forecasts are produced monthly, the MSL is

updated quarterly using these monthly forecasts.

Included in the monthly forecast is the Demand Deviation Report, which lists, by stock number, any changes from the previous forecast. This report requires a manual review by the OSPA and site representative, which enables them to investigate changes in the forecasted quantity. In some cases, adjustments to the forecast or historical demand data may be necessary. These adjustments require annotations to the report justifying the change. The OSPA provides the historical demand data, while the site representatives provide local insights into demand variation. These insights include knowledge about operational schedules, pre-planned usage for preventive maintenance, or the upcoming replacement of a weapon system.

Each quarter, the monthly forecasts are used to set the recommended MSL by NSN for each plant and a Level Analysis Report is also created. This report lists, by stock number, the previous order-up-to inventory level, the recommended MSL, the difference, and the unit price of the item for review. Using both the Demand Deviation Report and the Level Analysis Report, the site representative provides feedback on recommended level changes. The WSS OSPA considers the recommendations from the site representative and can manually adjust the inventory level setting.

Since the inventory levels are determined by statistical models, using historical demand data, and human input, the question arises as to how much of the forecast depends on the historical demand models and how much on human intervention. To illustrate this distinction, during the quarter level setting, site representatives recommended a change to the MSL generated by the model for only one of the thousands of line items. The first Level Analysis Report was generated in November 2012. In addition to the OSPA's review, senior leadership also examined the report before setting an inventory level. Again, only one of the thousands of line items drew the attention of senior leadership. The model recommended increasing the inventory level from 3 to 19. The unit cost of the item was over \$100,000. After verifying there were no issues in the level setting software, senior leadership asked the OSPAs to review the historical demand to justify the new level.

Setting the inventory level is also affected by the fixed allowance quantity assigned to some items. A fixed allowance is established prior to demand processing and applies to critical components necessary for a platform to perform its operational mission. The fixed allowance quantity is the minimum value to which an item's maximum stock level may be set, regardless of recommendations from the statistical model or the site representative. If the software application

calculates an inventory level below the fixed allowance quantity, based on demand data, then the level is set to the fixed allowance quantity. The inventory level established during the demand planning process is called the MSL. The reorder policy places an order up to the MSL value when the inventory level falls below the ROP.

The ROP is set to half the MSL in the Navy ERP system. For example, an item with an MSL of 10 will have an ROP of 5. If during the review period the inventory level is less than 5, an order is placed that will raise the inventory level back to 10. The number of items ordered is the difference between the MSL and the on-hand balance. The ROP calculation is simplistic, partly because of the desire to meet implementation milestones. Moving forward, the WSS plans to implement an algorithm to calculate the ROP based on the MSL, the average demand rate, and the time between consecutive inventory replenishments. This proposed algorithm is similar to how the legacy application calculated the ROP value.

While demand planning sets the MSL and ROP, supply planning is the process of replenishing the SNI based on these values. The Navy manages three groups of inventory, with different replenishment procedures for each. Since this analysis focuses on one of these groups, the replenishment procedures for this type of inventory are discussed. This particular group consists of the Defense Logistic Agency (DLA) and General Service Administration (GSA) inventory funded by the Budget Project 28 (BP 28). Table 1.1 lists the six budget projects that make up the Navy Working Capital Fund (NWCF). At the time of writing, balancing between sites does not occur, and the WSS is awaiting the results of this study before implementing balancing. The proposed BP 28 inventory replenishment process occurs in two cycles. The first cycle, Deployment Inventory Balancing, screens plants for inventory above the MSL that can be redistributed to plants with deficiencies. After redistributing the inventory between sites, the second cycle, Distributed Requirement Planning, places orders to fulfill the remaining inventory deficits. The replenishment process takes place twice a week at the NASs and FLCs, and daily at the TRFs.

The first cycle in the replenishment process requires business rules to govern the redistribution of inventory to avoid churn. Churn is the continuous movement of inventory between locations. Churn is undesirable because inventory in transit cannot be issued and the transportation cost may exceed the cost savings from using the excess inventory. The BP 28 inventory is held at 39 different locations around the world. The business rules that govern this process consider the unit price and the geographic locations of the inventory. Table 1.2 lists the balancing regions and

Table 1.1: Navy Supply Management is divided into six budget projects to organize the financial operations of the NWCF. (From Office of the Assistant Secretary of the Navy 2011)

	<i>Budget Project</i>
Wholesale	
Aviation Consumables	BP 34
Ship Reparables and Consumables	BP 81
Aviation Reparables	BP 85
Retail	
Ship's store	BP 21
General Consumables*	BP 28
Operations	
Operations and Reimbursables	BP 91

* Note: all inventory within this study is BP 28

the dollar value required to redistribute inventory between locations. For example, the Patuxent River NAS excess inventory with a unit price of \$69.00 could be redistributed to the Oceana NAS, but could not be distributed to Lemoore NAS, because Lemoore and Patuxent River are not in the same region. However, if the unit price of the inventory were \$76.00, then the inventory could be distributed to either site, since all sites being considered are contiguous United States sites. If Lemoore had a greater shortage than Oceana, the excess from Patuxent River would first be used to fulfill Lemoore's shortage, even though Oceana is closer. Then, any remaining excess would be transferred to Oceana.

Table 1.2: Inventory balancing regions are set up geographically, functionally (TRF), and are unit price dependent. (After Naval Supply Systems Command WSS 2011)

Region	Unit Price
Global (all sites)	\$5000 and above
CONUS (as well as Japan and Guam as own region)	\$75.00 to \$4999.99
Regional (NE, SE, NW, SW, Trident, Japan Less Guam)	\$25.00 to \$74.99
No balancing	Less than \$25.00

The balancing thresholds and regions in Table 1.2 are arbitrary. The official balancing rules and regions do not have any analytical support, since they were inserted as a placeholder to meet Navy ERP development milestones. This thesis will analyze these business rules to determine their viability. It will also explore alternative business rules and compare the outcomes.

Like the balancing regions and the dollar cost bands, the algorithm that determines the transfer

of inventory is also a proxy not supported by any analysis. This arbitrary inventory balancing algorithm uses the following process to determine which location with an excess is to transfer stock to a location with a shortage.

1. The system looks for the location with the (next) largest excess.
2. The system looks for the location with the (next) largest shortage.
3. The system checks whether both locations are in the same inventory balancing area. If this criterion is not met, the system returns to step 2 and determines the location with the next largest shortage, and so on. If the system does not find a permissible combination of locations, it returns to step 1, and determines the location with the next largest excess to check if a permissible combination exists. The process repeats. If the inventory balancing service does not find a permissible combination, inventory balancing does not take place.

The second cycle in the replenishment process fulfills the remaining deficiencies if the inventory level is below the ROP. The system generates a purchase request (PR) per plant for any deficiencies remaining once inventory balancing is complete. All PRs are reviewed by the WSS OSPA to ensure funding is available to purchase the inventory. The OSPA provides the replenishment package to the site representative for review. During the review, the site representative annotates recommendations to the replenishment package based on upcoming site requirements. The SPR returns the document to the OSPA for final approval. The software application obligates the approved PR and generates the purchase order (PO).

The supply planning process is representative of a more general inventory policy. Inventory policies are distinguished by the following two factors: the way replenishment orders are initiated and the decision rule that determines the size of the order. The following combinations of decision variables define different inventory policies: s (reorder point); R (review interval or order cycle); Q (order quantity); and S (order level). Common inventory policies are as follows: (s, Q) policy; (R, S) policy; and (s, S) policy (Silver et al. 1998, Chap. 7 and 8).

BP 28 replenishment within ERP uses a combination of (s, S) and (R, S) policies. With an (s, S) inventory policy, the reorder point s determines when an order is initiated. Unlike an (s, Q) inventory policy, which uses a fixed order quantity, the order quantity is the amount needed to bring the inventory back up to level S . An (s, S) policy assumes a continuous review (i.e., the inventory level is tracked continuously). However, supply processing in the ERP system is not continuous, but is reviewed daily for the TRFs and twice weekly for the NASs. In this case, the inventory policy is characterized by a third parameter that specifies the length of the review

interval, R . Using the same notation, this is the (R,s,S) policy. In the case of the Navy ERP system, R is daily or weekly, depending on the site, s is the reorder point value, and S is the maximum stock level value.

1.2 Literature Review

To establish a theoretical framework for inventory balancing within Navy ERP and to evaluate alternative balancing policies, a review of previous studies and literature helps define terminology and provides a common method to examine the issues. Inventory balancing is a special type of lateral transshipment. The term transshipment refers to sharing inventory among locations while the word lateral signifies that the sharing occurs at the same echelon within the supply chain. A lateral transshipment policy is a method to minimize inventory levels while maintaining a given service level. (Silver et al. 1998)

Although a policy definition of lateral transshipment specifies service level as a performance metric, this study uses fill rate instead since the Navy uses this metric of performance. Whereas, a fill rate is the percentage of demand that is met on time, service level is the probability that the demand is satisfied while reorders remain outstanding. In the long run, this probability should be the percentage of order cycles without a backorder. Both fill rate and service level are measured by percentages. However, for the same percentage, the service level metric requires considerably more safety stock than the fill rate metric.

For example, if the service level is set to 0.9, this means there is a 90 percent chance that the demand is met prior to receiving reorders. Suppose the reorder frequency is 10 times per year. This means, on average, 9 of those 10 order cycles won't have a backorder but allows for more than one shortage in that cycle. To meet this service level, a retailer must carry enough safety stock so that they don't have any shortages in 9 order cycles. So, a service level set to 0.9 means carrying a high safety stock and having very few unsatisfied customers. By contrast, if the fill rate is set to 0.9, this results in dissatisfying 10 percent of the customer orders, which requires less safety stock than compared to a service level set to 0.9.

According to Silver et al. (1998) supply chain management literature divides inventory decisions into two types: structural and coordination. While structural decisions are often long lasting and infrequent due to the associated high costs, such as how many retail sites to have and where to

locate them, coordination decisions regularly occur once the structural decisions are in place. Variation in lead times and demand play a significant role in coordination decisions. This study focuses on coordination decisions, primarily to identify the best transshipment policy without changing the structure of Navy ERP.

Before addressing decisions associated with lateral transshipment policies, one must first consider the distribution of information and where control lies in the organization. The supporting literature articulates this as local versus global information and centralized versus decentralized control. According to Silver et al. (1998), the use of global information and centralized control produces better results since the decision maker's focus is on overall system performance using all available information. This leads to more accurate and efficient supply chain decisions that can greatly minimize large replenishment orders at individual retail sites. However, most firms fail because of the difficulty to provide the cooperation and coordination across the organization in order to achieve success.

The Navy tried and failed to institute the SNI due to the difficulties of sharing information over multiple databases using legacy systems. Conversely, Navy ERP shifts the enterprise towards global information and centralizes control of the system, allowing decision makers to have visibility of the demand, costs, and inventory status across all locations. These coordination challenges across the organization are not unique to the Navy.

In fact, the complexity of this problem is only magnified for competitive firms that are hesitant to share sensitive data with their suppliers, much less their competitors. Independently owned retail sites have a greater incentive to optimize their own profits rather than the entire system. For these reasons, very few firms operate under centralized control while sharing information globally. Therefore, literature regarding the Navy's situation is scarce since coordinating across multiple operations, with different functions, and sometimes between firms, is very challenging and rarely implemented. Rather, the majority of the literature available emphasizes the competitive nature among retailers to find Nash equilibrium, as Cachon (2001) demonstrates. Cachon's (2001) research studies competitive firms and finds, under some circumstances, the Nash equilibrium is equal to the cooperative reorder policy, which means the competitive solution is not necessarily inefficient. However, under other circumstances, he finds the competitive costs at the Nash equilibrium are much higher than the cooperative solution. To ensure the supply chain operates at optimal reorder policy, Cachon (2001) shows that competitive firms can reach the cooperative solution by having the supplier set all the reorder points.

Since most literature concentrates on competitive firms, it often assumes transshipment among retailers is not allowed. Rudi, Kapur, and Pyke's (2001) study is an exception that looks at transshipment between competitive firms. Their analysis, like Cachon (2001) looks to find ways that competitive firms can reach the cooperative reorder policy. Unlike Cachon (2001), who recommended giving reorder control to the supplier, their research finds a transshipment price that facilitates reorder points similar to the cooperative solution.

However, even cooperative retailers must consider if a transshipment policy benefits their enterprise. This study examines whether to implement a lateral transshipment policy and the various policy choices. Navy ERP's Deployment Inventory Balancing is a preventive lateral transshipment process. It is preventive because it prevents a future shortage. Allen (1958) was one of the pioneers studying how to redistribute stock among several locations to minimize shortage, holding, and transportation (transshipment) costs. The Bureau of Supplies and Accounts of the U.S. Navy sponsored his research to improve inventory efficiency. Tagaras (1999) refers to models, such as Allen's (1958), that minimize future risk of shortages as preventive transshipment, as opposed to those that prevent an impending stockout, which he calls emergency transshipment. Emergency transshipment is moving materiel from a retailer with abundant stock to a retailer facing a stockout from current demand, while preventive transshipment is redistributing stock to mitigate the risk of future stockouts before demand occurs.

Perhaps Jönsson and Silver's (1987) study most closely resembles the lateral transshipment challenges facing the Navy. Their study considered a central warehouse and multiple retail sites. Similar to the OSPA located at WSS Philadelphia, the central warehouse makes decisions for all of the sites and manages the replenishment process without retaining any stock. However, a major difference in Jönsson and Silver's (1987) study is that the central warehouse acts as the distribution center to the retail sites. Whereas with Navy ERP, replenishments ship directly from the supplier to the site. The study also differs in that Jönsson and Silver rebalance inventory between sites in the middle of the reorder cycle, while the redistribution in Navy ERP occurs immediately prior to placing a replenishment order. Unlike Navy ERP, Jönsson and Silver (1987) also assume that the demand per site has the same probability density function.

Acknowledging the differences noted previously, Jönsson and Silver's (1987) findings provide a general context for when a rebalancing policy may be beneficial and when it is not. In their study, they compare balancing across all sites regardless of transportation or unit cost, i.e., a complete redistribution policy, to not redistributing among sites. For a given service level, the

redistribution policy requires a lower order-up-to value S than a policy without redistribution, resulting in a smaller inventory investment. However, Jönsson and Silver (1987) show there is a trade off between holding costs and shipment costs, which determines the viability of a redistribution policy. Their study shows that a redistribution policy becomes more beneficial as demand increases, the number of retail sites increase, lead-time (Both between retail sites and from the supplier to the sites) becomes shorter, and the desired fill rate increases.

The Navy's transshipment policy does not take advantage of increasing demand, as in Jönsson and Silver's (1987) results, since it is based solely on unit price and geographic distance between sites. This allows some low demand, high unit price transshipments and prevents high demand, low unit price transshipments. This study will consider alternative transshipment policies that take into account the benefits of Jönsson and Silver's (1987) findings.

The benefits found by Jönsson and Silver's (1987) study are contingent upon the inventory savings outweighing the transshipment costs to meet mandated fill rates. To quantify this benefit, one must calculate transshipment and holding costs. Calculating shipping costs is relatively straightforward. However, holding costs are more difficult to compute since the cost of keeping materiel in inventory is comprised of multiple expenses. These include the cost of capital, warehousing expenses, labor costs, deterioration of stock, and obsolescence. Silver et al. (1998) offer this common formula to calculate holding costs

$$\text{Holding costs per year} = \bar{I}vr \quad (1.1)$$

where $\bar{I}$ is the average inventory in units, v is the per unit price, and r is the carrying charge of one dollar of inventory per year. Of the three variables, r is the most difficult to compute. It incorporates both the opportunity cost, the expected rate of return ceded by bypassing other potential investments, and the cost of storage, which depends on bulkiness and special handling instructions. For example, storage costs at the Navy retail sites are based on extended cubage ($\text{qty} \times \text{item cubage}$).

However, in comparison to these storage costs, the more difficult calculation is the opportunity costs associated with tying up capital. This difficulty arises from the myriad of alternative investment options available and the propensity for such investments' return to change day-to-day. The U.S. Navy, like all firms, wrestles with this challenge. Spackman (2004) considers the cost of capital in the public sector and finds it is fundamentally different than the private

Table 1.3: The following are the fiscal year 2013 rates for DLA distribution. (After McNeeley 2013)

	Annual ($\frac{\text{dollars}}{\text{Cubage}}$)	Monthly ($\frac{\text{dollars}}{\text{Cubage}}$)
Covered	5.48	0.4567
Open	0.53	0.0442
Specializes	7.61	0.6342

sector. This is because the cost of capital in the public sector must incorporate the tax paid by the private sector and must also account for the cost of investment risk. Spackman (2004) states that politics perpetuates the simple-to-explain, however flawed, conception that opportunity cost is the cost of relinquished private investment. He recognizes within government two opposing forces: comptrollers who support overestimating the cost of capital and executors who support underestimating it, Comptrollers believe higher estimation will restrain expenditures, while executors see more assets as a way to mitigate risk. This study uses the carrying charge, r , from previous studies to calculate the cost of capital. The last official value of r on record is in a GAO 1995 that finds that r varies by site around 22 percent. However, WSS adopted a much lower value of r equal to 10 percent. This study will assume the 10 percent holding cost rate in lieu of a more recent value.

Jönsson and Silver's (1987) study shows under what conditions a redistribution policy is beneficial, but given the differences in Navy ERP, not all of the study's assumptions might be met by the Navy. If Jönsson and Silver's (1987) assumptions are not met, then the criteria of when a redistribution policy is beneficial may change. McGavin, Ward, and Schwarz (1997) deviate from Jönsson and Silver's (1987) assumption that the demand between sites is identical and normally distributed, assuming, instead, that distributions are asymmetric. With this assumption, balancing is not necessarily optimal. Allen (1958) also does not assume identical demand distributions; however, his results consider only a single-time interval. When multiple intervals are considered, McGavin et al. (1997) have shown that you cannot always assume balancing provides an optimal solution. Optimality fails in cases where a central warehouse lacks the materiel to bring all sites to the order-up-to level. However, if there is sufficient materiel, McGavin et al. (1997) find balancing is optimal, even with asymmetric demand distributions, as long as backorders are allowed. In Navy ERP, the system never lacks materiel to bring all sites to the order-up-to level, since each site orders directly from suppliers, not from a central warehouse with minimal stock available. As long as the suppliers can fulfill the purchase orders, then the allocation (balance) assumption is met.

The models discussed so far have been only analytical, and they make several simplifying assumptions, in part, to provide tractable solutions. However, practical applications of transshipment may be more complex, and an analytically tractable solution may be infeasible. Several researchers have used simulations to build upon analytic models in an effort to incorporate realism in their analyses. At times, the complexity of real-world lateral transshipment problems makes finding an analytical solution extremely difficult. The strength of simulation models lies within their ability to answer these complicated problems unsolvable by analytic means. For this reason, Ekren and Heragu's (2008) study chose a simulation-based optimization procedure to evaluate inventory balancing policies in order to maintain the complexity of the system. There is a trend in literature to expand upon analytic models by relaxing assumptions through simulation modeling. This relaxation makes formulating a closed form solution difficult. In most cases, the insights from the analytical model hold with the relaxation of assumptions in the simulation.

Both analytic and simulation researchers select different distributions to model demand and lead time. Analytical studies often use a convenient distribution to make their models tractable, disregarding if the distribution is realistic. Whereas, simulations can use almost any distribution since tractability is not an issue. Jönsson and Silver (1987) use the normal distribution in their analytical model, in part, because its properties are well known. Tyworth and O'Neill (1997) devote an entire paper on evaluating the misspecification error from assuming lead-time demand is normally distributed. They use historical demand, not normally distributed, from various industries to test the sensitivity of models that assume demand is normal. Tyworth and O'Neill (1997) choose industries whose demand distribution is skewed left, right, and not at all. Their results provide evidence that the normal approximation is robust. They find that the error associated with misspecifying the normal distribution shrinks as target fill rate decreases from 99 to 80 percent.

Although the normal approximation is robust, other authors use the Poisson. The Poisson is a discrete, non-negative distribution that is appropriate to depict the probability of "rare" events. For example, if demand arises from a failure process, as in repair parts, there is a theoretical justification for the Poisson. However, Needham and Evers (1998) do not use the Poisson distribution to model demand because it requires the mean and variance to be equal. Instead, they use a truncated version of the normal distribution to prevent negative values and to circumvent the Poisson requirement of equal mean and variance. These distributions are less useful, however, when the item is expensive and demand is sporadic and slow moving. In this case, Kukreja and Schmidt (2005) find that a Weibull distribution performs better. Whereas,

Wong, van Houtum, Cattrysse, and Van Oudheusden (2006) use the Poisson process to describe the repair part distribution for airplanes. They argue that Poisson is well suited when machinery is in constant use, high system availability is a major consideration, and keeping inventory costs as low as possible. Regardless of distribution used, simulations are able to incorporate different distributions and even compare and contrast scenarios with varying underlying distributions.

Banks, Buckley, Jain, Lendermann, and Manivannan's (2002) study notes the following issues with supply chain management systems analytic approach to inventory policy decisions immediately prior to implementing them: the inability to change parameters in the past, the disruptive or inability to experiment with the real system, the inability or difficulty to represent random effects, and the difficulty to see the effect of changes in the long run. Their research addresses how simulation solutions have been effectively used in other studies to address the previously listed shortcomings. Banks et al. (2002) distinguish between two simulation approaches, the "top-down-approach," in which a simulation is built from scratch and modified over time to provide more detailed analysis, and the "bottom-up approach," in which a pre-existing complex simulation model is modified to suit a problem's particular needs. This thesis uses the first type, building the simulation from scratch using Excel and Visual Basic for Applications (VBA) to simulate the Navy BP 28 supply chain.

Amaral and Kuettner (2008) use Excel and Visual Basic for Applications (VBA) to illustrate that using only these two tools can improve the supply chain decision-making, as evidenced by Hewlett-Packard's operations research department. Amaral and Kuettner, former employees of Hewlett-Packard, find that spreadsheet modeling is particularly well suited to inform supply inventory policy decisions. They argue that a precise representation of a complex supply chain is not needed to answer decision makers' questions and the data to attempt such a model probably does not exist. Instead, they find that spreadsheets force modelers to simplify their assumptions and that the final model is easier for the decision maker to understand. They find this is because a spreadsheet is less intimidating to non-technical decision makers than other analytical techniques.

In addition to the supporting literature on lateral transshipment policies, determining the reorder point and order-up-to-level is an area of study that requires investigation to support this thesis. Changes in these values over time create the excess and shortages that make redistribution of inventory feasible. Although this paper focuses on preventive lateral transshipments to balance inventories across sites, such imbalance arises as a result of level setting. Variation in demand is ultimately the cause of resetting of the ROP and MSL. Unfortunately, the algorithm that

determines the maximum stock level (MSL) in Navy ERP is proprietary, and, therefore, the calculation that sets the inventory level is unavailable for this study. However, the legacy is available, but it depends on lookup tables no longer tied to the current stock.

This study uses a combination of the legacy code and recent literature to create a hybrid algorithm, which determines the reorder point and order-up-to-level (or MSL) for a given fill rate. Most of the literature, however, uses cycle service level to determine the reorder point and order-up-to-level and focuses on other more tractable inventory policies than the Navy's (R, s, S) inventory policy. Fortunately, Silver, Naseraldin, and Bischak (2008) provide a simple method that uses a target fill rate to calculate s and S . Their method addresses the delay between when the inventory level breaches the reorder point and when the subsequent order is placed within the (R, s, S) inventory policy. Silver et al. (2008) assume demand is normally distributed and lead time is constant; however, simulation supports that other distributions (Poisson) will provide reasonable results as long as coefficient of variance (CV) is less than .5, since Silver et al.'s (2008) process uses a normal distribution, values greater than 0.5 for the CV could result in negative values. To avoid negative values a gamma distribution would likely be more appropriate.

The CV measures the stability of a item's demand, comparing the variability in demand to the magnitude of the average demand. The ability to accurately forecast demand with high variability is difficult, if not impossible. Therefore, high levels of inventory must be held to meet even minimal fill rate requirements. In contrast, lower variability means that demand forecasts are much easier to accurately forecast.

$$CV = \frac{\sigma_X}{E(X)} \quad (1.2)$$

where

$$\begin{aligned} \sigma_X &= \text{standard deviation of demand} \\ E(X) &= \text{expected value of demand} \end{aligned}$$

To test the robustness of their method, simulation runs using a compound Poisson distributed demand and variable lead time confirms that the normality assumption can be relaxed. Simulating different demand distributions, the authors' method deviates only slightly from the target fill rate. The simulation in this study uses Silver et al.'s (2008) algorithm to calculate s and S .

1.3 Purpose of the Analysis

With the creation of an SNI, the Navy ERP system provides unprecedented visibility of Navy-owned inventory at onshore retail locations. This visibility enables the redistribution of excess inventory to locations with shortages to avoid unnecessary procurement of GSA and DLA inventory. The goal of this thesis is to assess inventory balancing business rules, including the option of not balancing, their effect on total cost, and their ability to meet demand.

To do this, a computer simulation provides the expected per order cycle, replenishment, stockout, inventory, backorder, and transshipment costs. The simulations calculate these costs per period using a mathematical model. Together, these costs make up the total cost for a given set of business rules. Analyzing both the expected total cost and its individual components (holding, backorder, replenishment, and transshipment costs) will provide one mechanism in which to compare the cost efficiencies gained or lost using a given set of inventory balancing business rules. In addition to cost, the simulation also provides the expected fill rate, which is a measure of how effectively inventory meets demand. It is important to emphasize that this study focuses on the long-term costs realized over a two-year period. These long-term costs may misrepresent the short-term costs immediately following implementation because of the build-up of imbalances within the system. However, the longer the balancing policy is implemented, the more the initial imbalances should decrease.

1.4 Scope, Limitations, and Assumptions

The scope of the proposed analysis considers three locations of the 39 that hold BP 28 inventory, two within the same region and one outside. At the time of writing, redistribution between sites does not occur. Again, this study uses a computer simulation and a mathematical model to assess the impact of the business rules on the overall cost and supply effectiveness. Historical demand data is run through the “without redistribution” model to create a base case against which the merits of the simulation can be judged. Ideally, this base case should produce results that are consistent with the historical data. For a given period, the study compares no balancing, balancing under the proposed business rules, and balancing without unit price limitations based on the resulting expected total costs and fill rates. Based on the outcome of the comparison, modifications to the business rules are proposed, tested, and compared.

1.5 Research Questions

This study examines six research questions to determine the benefits of competing transshipment policies:

1. Do fill rates at individual retail sites improve when implementing a preventive lateral transshipment policy (inventory balancing)?
2. Given fill rates decrease, how do the different demand distributions affect the benefit of the preventive transshipment policy?
3. Regardless of whether or not there is an improvement, how does redistribution affect procurements?
4. Do different balancing criteria at the retail level result in differing fill rate between the retail establishments? If the proposed balancing criteria fail to improve on those policies with no balancing, do alternative criteria (i.e., a different set of balancing business rules) provide a benefit?
5. If any of the sets of balancing decision rules tend to improve fill rates, does one set of balancing criteria tend to outperform the others? If so, which of the alternatives reduce procurement costs by more than they increase transportation costs?
6. Assess which policies tend to maintain lower inventories over time. What is the impact of preventive transshipment on the amount of excess and stock above the MSL as compared to the case with no inventory balancing?

CHAPTER 2:

Data and Methodology

This study determines and compares the total costs and fill rates of competing inventory balancing policies. Decision makers can use this information to evaluate each policy prior to implementation to determine which policy best meets their needs. Several interviews conducted during a one-week site visit to Mechanicsburg and Philadelphia produced the data and influenced the methodology used in this report. We collected information on inventory procedures and areas of concern for the Weapon Systems Support (WSS) and Naval Supply Systems Command (NAVSUP) headquarters by interviewing military and civilian experts. The site visit not only provided electronic data necessary for this study, but also the context for the study. Continued conversations and communications in the months following the site visit provided additional data and information that helped refine and enhance the modeling effort.

In Section 2.1, we examine Budget Project 28 (BP 28) inventory data for Lemoore, Oceana, and Patuxent River to determine whether these sites show potential for inventory balancing. BP 28 inventory comprises consumable items, such as aircraft arresting gear, hoses, clamps, and valves, and is the only Navy-managed inventory eligible for inventory balancing. Our approach begins by assessing the inventory excess and shortage at each location. Next, we calculate the overlap in demand among the three locations and determine which items have sufficient overlap to be included in this study. Finally, we consider the demand pattern of the selected items over two years of data.

Careful examination of historical data can show if demand overlaps exist and can help identify items to study. However, no historical balancing data exist to measure the potential cost savings of rebalancing. Therefore, we construct a model to estimate these cost savings. The model computes the total relevant cost per day for each item. The total cost is composed of holding, backorder, procurement, and transshipment costs. The backorder and transshipment costs include a shipping cost component. We develop a shipping cost model in Section 2.2.1 and describe the inventory management cost model in Section 2.2.2. Analytic solutions rely on either constant or convenient distributions of demand and lead times. These assumptions simplify the inventory analysis to make the problem tractable. However, they do not reflect the observed demand and lead time data, so we develop a simulation model to estimate the costs in Section 2.3.

2.1 Data

The Navy owns hundreds of millions of dollars in excess inventory, while at the same time facing stock shortages. Excess inventory is defined as the number of units in stock above the maximum stock level (MSL), or order-up-to level (denoted by S), while a shortage is defined as the amount of inventory below the MSL. Finding better business rules for inventory balancing is a step towards efficiently using the inventory on-hand, potentially realizing cost savings. In Section 2.1, we analyze the potential for an inventory balancing transshipment policy, given the inventory data from WSS for the naval air stations (NASs) at Lemoore, Oceana, and Patuxent River. We do this in Subsection 2.1.1 by first stratifying the data by the amount of excess and shortage each site carries. From this stratification, we identify the quantity of excess items that fall within each of the WSS proposed unit price thresholds. The thresholds use unit price to determine eligibility for transshipment between geographical regions. Further analysis in Subsection 2.1.2 then examines the commonality of demand among the three sites. After identifying items with common demand, in Section 2.1.3, we analyze those items with the largest quarterly dollar volume and look at their historical demand pattern over two years. This focuses our research on inventory that has the largest impact on inventory management costs.

2.1.1 Excess and Shortage

For inventory balancing to occur, the system must have shortages and excesses to balance. Of the two, excess is the more important parameter, as balancing cannot occur without it. Furthermore, future demand will generate shortages in the system, even if there are none at present. Figure 2.1 illustrates how much inventory is above, below, or equal to the MSL for the three locations.

The data in Figure 2.1 is drawn from an Navy Enterprise Resource Planning (ERP) *Excess and Deficit Detail Report* run on October 16, 2012, for Naval Air Stations Lemoore, Oceana, and Patuxent River. For Lemoore, Oceana, and Patuxent River, the percentage of items in stock above or without an MSL are 19 percent, 27 percent, and 7 percent, respectively. Table 2.1 lists the value of excess inventory by location and the percentage of the inventory value held by each location.

After establishing the excess held at the three locations, we then examine how this excess is distributed into the WSS proposed unit price thresholds, as listed in Table 1.2 in Chapter 1. For example, excess inventory less than \$25 will be ineligible for balancing under the WSS proposed business rules.

Excess and Shortage Material Status




Figure 2.1: This bar chart shows the number of inventory line items each location holds below, equal to, and above the MSL. (After Cardillo 2012)

Table 2.1: Excess inventory value and percentage of inventory. (After Cardillo 2012)

<i>Location</i>	<i>Excess(\$)</i>	<i>Inventory(\$)</i>	<i>Percentage(%)</i>
Lemoore	\$24.6M	\$1.7B	1.4%
Oceana	\$9.7M	\$745M	1.3%
Patuxent River	\$9.1M	\$226M	4.0%
Total	43.4M	2.7B	1.6%

Figure 2.2 displays a breakdown of the excess inventory by number of line items and by overall dollar value of the items within each of the dollar thresholds. Over half the excess inventory is ineligible for transshipment under the WSS proposed dollar thresholds, as shown in Subfigure 2.2a. While this will probably minimize churn by limiting the number of line items for transshipment, it also limits the potential decrease in procurements. This study considers the impact on overall costs by removing the dollar cost thresholds. Conversely, Subfigure 2.2b shows that over three-quarters of the excess inventory dollar value falls in the highest unit price dollar thresholds. The two subfigures reflect an inverse relationship in that the largest dollar threshold by quantity is the smallest in terms of dollar value.

Comparing the two subfigures demonstrates that the majority of excess line items fall in the “less than \$25.00” dollar cost threshold. In contrast, the majority of the excess inventory by dollar



Figure 2.2: The two subfigures contrast the difference between the number of line items and the dollar value that are within each of the WSS proposed unit price dollar thresholds, which determines the eligibility of the item for transshipment. Although 72 percent of all the line items fall within the \$25.00 unit price threshold, 82 percent of the total dollar value of items is in excess of \$5,000.

value falls within the “above \$5,000.00” category. This indicates potentially high savings using transshipments, with limited churn.

2.1.2 Commonality

In addition to establishing the amount of excess and shortage each site carries, an inventory balancing policy requires that each location has demand for common parts to facilitate transshipments among the three sites. If each location’s demand is unique, then no transshipments are possible. The following analysis examines the commonality in demand for line items at each location.

Figure 2.3 displays the commonality of parts between the three locations using a Venn diagram. The overlapping circles represent the demand for individual line items at each of the three locations. Each circle represents a different location. The sum of the numbers within a circle represents the demand for unique items at that location, while the numbers outside that circle represent inventory items not held at that location. For instance, in this three-set Venn diagram, there are circles representing the inventory items held at each of Lemoore, Oceana, and Patuxent River. The overlapping areas, or intersections, represent the common demand for unique items at each of the locations. The intersection of all three circles shows that there are 45 items in demand at all three locations. Note that Figure 2.3 is not drawn to scale.

From Figure 2.3, 809 line items have demand common to at least two of the three locations. Table 2.2 shows the number of common line items that fall into each of the unit price thresholds. As shown in Table 2.2, only two common items are over the \$5,000 unit price threshold, and the




Figure 2.3: The following Venn diagram shows commonality of demand shared among Lemoore, Oceana, and Patuxent River. The circles are not drawn to scale. (After Liskow 2013)

majority are less than \$25.00 in value.

Table 2.2: Common line items that fall within each of the WSS proposed unit price thresholds. The majority of common items are less than the \$25.00 minimum needed for transshipment eligibility under the WSS proposed balancing rules. (After Liskow 2013)

	Lemoore	Oceana	Patuxent River
\$ 5,000 and above	15	16	3
\$ 75.00 to \$ 4,999.99	135	218	100
\$ 25.00 to \$ 74.99	43	84	55
< \$ 25.00	187	456	351
Total	380	774	509

* Note: 45 items are carried at all three locations

2.1.3 Unit Price versus Dollar Volume

Silver et al. (1998) advocate inventory decisions based on the contribution of an individual item to the annual dollar volume. In their extensive study of inventory systems, they found that, in general, 20 percent of items in an inventory account for 80 percent of the total annual dollar volume (Silver et al. 1998). Figure 2.4 shows that 10 percent of the National (or NATO) Item Identification Numbers (NIINs) in demand at Patuxent River account for 78 percent of their

total quarterly dollar volume, while 10 percent of the NIINs in demand at Lemoore account for 86 percent of its total quarterly dollar volume. Over 96 percent of Oceana's dollar volume is reflected in only 10 percent of their NIINs.




Figure 2.4: Distribution by Value of NIINs.

With so much of the overall value of inventory dependent on only 10 percent of the items, the decisions we make concerning this inventory can make a significant difference. Therefore, inventory balancing can achieve significant improvements in costs and fill rates simply by making this 10 percent eligible for balancing. The WSS proposed business rules use unit price thresholds to regulate eligibility for balancing. However, this approach ignores items that fall within this 10 percent and includes items outside of this grouping. In addition, some of the items within the 10 percent have a unit price of less than a dollar, while some outside this 10 percent exceed a thousand dollars. Therefore, unit price alone is a poor indicator of the significance of a particular item to inventory. By stratifying items by dollar volume, we can identify critical items that have the most impact on inventory management costs. The rest of the inventory contributes little to overall costs.

To create Figure 2.4, we first multiply the unit price by the quarterly average demand. We then arrange the resulting dollar volume from greatest to smallest, as seen in Table 2.3 for Patuxent River's inventory. Once the items are arranged in descending order, we assign an item number corresponding to where it falls in the list. For example, the first item in Table 2.3 has the highest dollar volume of the items in demand at Patuxent River. The third column in Table 2.3 is calculated by dividing the item number by the total number of items, and the sixth column

is calculated by dividing the cumulative dollar volume by the total cumulative dollar volume. Plotting the third column against the sixth column yields the plots in Figure 2.4. Table 2.3 is useful for identifying items that have the greatest potential to influence overall costs.

Table 2.3: Patuxent River listing of NIINs by descending dollar demand. (After Liskow 2013)

Item Number	NIIN	Cumulative Percent of NIIN	Quarterly Dollar Volume	Cumulative Dollar Volume	Cumulative Percent of Total Volume
1	015290122	0.021	285,187.00	285,187.00	10.35
2	015569064	0.043	131,925.50	417,112.50	15.131
3	014732035	0.064	40,461.43	457,573.93	16.599
⋮					
4678	002483835	99.978	0.053	2,756,705.29	99.99
4679	000043134	100.00	0.045	2,756,705.29	100.00

This study examines five common items to evaluate the inventory balancing transshipments policies. Table 2.4 lists the five selected items. These items are chosen based on having the highest dollar volume per quarter, which is calculated by multiplying the total demand per quarter by the unit price per item. Two of the resulting five items are less than \$25.00, so are ineligible for transshipment under the WSS proposed balancing policy.

Table 2.4: Summary data for the five largest common items by quarterly dollar volume. (After Liskow 2013)

NIIN	Weight (lbs)	Cube (ft ³)	Nomenclature	Quarterly Dollar Volume		
				Unit Price	Qty	Total
011506504	37	0.133	VALVE, REGULATING, TEMPERATURE	1154.61	14844.5	17139608.15
011669450	0.9	0.008	WIRE ROPE ASSEMBLY, SINGLE LEG	81.23	169.5	13768.485
001849417	0.1	0.003	SEAL, BOSS	38.4	339.25	13027.2
014617380	10	0.1481	HOOK POINT, ARRESTING	8676.59	0.75	6507.4425
001003144	0.4	0.008	CONE AND ROLLERS, TAPERED ROLLER	11.34	546.75	6200.145

* Quarterly dollar value is the highest value of the three sites

2.1.4 Historical Demand

We next examine the demand patterns for the five chosen items. The following parameters are inputs into the simulation model to compare the balancing policies. *Demand* is a measure of the

customers' need for a particular item. *Lead time* is the amount of time between when a purchase order (PO) is submitted and when it is fulfilled. An *Allowance* is the minimum order-up-to level for an item. Although an item may not have a previously recorded demand or demand supports a lower MSL, allowances for stocking are based on other criteria. The allowance quantity is based on how essential a weapon system is for operational success rather than previously recorded demand. For each line item, the WSS provided the average quarterly demand, quarterly demand variance, the average lead time, inventory allowance, and the monthly demand total for 24 months.

Figure 2.5 represents the historical demand for the three locations. Subfigure 2.5a for NIIN 014617380 exhibits a pattern of very low or zero demand in a given month followed by spikes of demand in the following months. In other NIINs, the difference in demand between months can be 0 to 5,000. This pattern is termed lumpy demand. Here, forecasting becomes more difficult as the historical demand history exhibits an increasingly lumpy demand. The five items listed in Table 2.4 are components required when replacing an original part due to failure, age, or other reasons. These components have inherently lumpy demand, since their demand is driven by part failures, which occur intermittently. Another instance of lumpy demand occurs when organizations wait to make a mass purchase, possibly due to quarterly or annual budgets or the long life cycle of an item.

Since the WSS follows a (R, s, S) reorder policy, the service part failure argument best explains the lumpy demand in this case. However, the extremes in demand for some items between months make this explanation difficult to accept. The unit price for NIIN 011445679, which is not shown, is \$8.97. The demand for this item fluctuates between 0 and 5,664 each month. The lumpy demand in this instance might be due to the WSS purchasing in bulk when funding is available over the MSL. For example, the ERP system generates a purchase requests (PRs) to bring the inventory stock level up to the MSL. However, the Operational Supply Planning Analyst (OSPA) may increase the replenishment quantity 100 times above the generated PR to use a budget surplus, regardless of demand. In contrast, when funding is unavailable, the OSPA may not approve a PR for several months, even though demand may warrant it.

Our analysis requires a reasonable estimate of the demand, but the WSS does not report customers' demand. Instead, they report the warehouse's inventory replenishment orders, not the number of issues to the customer. Clearly, there is a connection between warehouse replenishment orders and actual customer issues. However, because of budget issues, as in the previous




Figure 2.5: The two subfigures contrast the differences in demand patterns at two locations for two items. Each subfigure shows the difference in demand patterns, demand by month for nine months, for the same item between Patuxent River and Oceana. Notice the demand patterns vary both between items and locations. (After Liskow 2013)

example, or for other reasons, warehouses may replenish in excess of their MSLs without supporting demand, or may be unable to replenish up to their MSL when there is demand. We use the replenishment order data to represent historical demand in our analysis, recognizing the shortcomings in doing so, and note that using the actual customer issue figures would better represent demand. One of the advantages of ERP is the visibility of data it provides across the enterprise. By not reporting the individual issues at each site, the WSS is limiting the benefits of this tool.

2.2 Model Formulation

In this section, we address the question of how to represent the total relevant costs in a model. A model provides a simplified representation of a system over time. Relevant costs include only those costs that impact inventory balancing decisions. To identify relevant costs, the model calculates the holding, backorder, procurement, and transshipment costs. Since the backorder and transshipment costs have a shipping cost component, they require a separate model. In Subsection 2.2.1, we develop a shipping cost model. In Section 2.2.2, the shipping cost model is integrated into the main model.

2.2.1 Two-Day Shipping Model

The two-day shipping model calculates the cost of shipping inventory between locations based on the weight of inventory shipped. This prediction model is built using observed shipping costs provided by the WSS. Figure 2.6 depicts the relationship between the weight of an item and its shipping cost.




Figure 2.6: A scatter plot of the shipping data. (After Palko 2013)

As shown in Figure 2.6, there is a positive relationship between weight on the x-axis and shipping cost on the y-axis. Linear regression finds the best-fitting straight line through the points, called a regression line, by minimizing the sum of the squares of the differences between the predicted cost and the observed cost over all observations.

The difference between the observed value and the estimated function value is

$$\Delta_i = y_i - \hat{y}_i \quad (2.1)$$

where

y_i = observed shipping cost

$\hat{y}_i$ = predicted shipping cost using a regression line

The formula for the predicted shipping cost is shown in Equation 2.2. Equation 2.2 provides a

cost estimate based on the weight of the item. The quality of this prediction depends on the value of the slope coefficient and the intercept value.

$$\hat{y}_i = \hat{\beta}_0 + \hat{\beta}_1 x_i \quad (2.2)$$

where

$$\begin{aligned} \hat{y}_i &= \text{predicted cost to ship to location } i \\ \hat{\beta}_0 &= \text{intercept value} \\ \hat{\beta}_1 &= \text{slope coefficient for weight} \\ x_i &= \text{the weight of the items shipped to location } i \end{aligned}$$

The values for $\hat{\beta}_0$ and $\hat{\beta}_1$ are calculated by minimizing the sum of the squared difference between the observed value and the predicted value across the sample observations. Optimizing Equation 2.3 for the given paired weight and cost data, $D = \{(x_1, y_1), \dots, (x_n, y_n)\}$, determines the value for the scalar coefficient, $\hat{\beta}_1$, and intercept value, $\hat{\beta}_0$. Equation 2.3 estimates the values of $\hat{\beta}_0$ and $\hat{\beta}_1$. See page 569-572 of Wackerly, Mendenhall, and Scheaffer (2008) for an example of how to solve for $\hat{\beta}_0$ and $\hat{\beta}_1$.

$$\min \sum_{i=0}^n \left(y_i - x_i \hat{\beta}_1 - \hat{\beta}_0 \right)^2 \quad (2.3)$$

Equation 2.4 simply uses all the data points to estimate the parameters for its linear model.

The resulting linear model is

$$\hat{y} = 3.65363 + 0.77097x \quad (2.4)$$

where

$$\begin{aligned} \hat{\beta}_0 &= 3.65363 \\ \hat{\beta}_1 &= 0.77097 \end{aligned}$$

From the observed data, there seems to be a minimum price of \$4.15 to ship inventory, regardless

of weight. Equation 2.5 uses only the observations above \$4.15 to estimate the parameters of the linear model. Here, the shipping costs never fall below \$4.15, regardless of weight. The resulting linear model is then set equal to \$4.15 and solved for x . This value represents the weight at which the model predicts \$4.15, in this case, 3.5 pounds. In Equation 2.5, if the weight is below 3.5 pounds, then the predicted cost is \$4.15, and if it is above 3.5 pounds, the linear model provides the predicted cost.

The resulting model is

$$\hat{y} = \max [1.197527 + 0.843028x, \$4.15] \quad (2.5)$$

Table 2.5 shows the sample data and the output from Equation 2.4 in the third column and the output from Equation 2.5 in the fourth column.

Table 2.5: Observed paired data for weight and cost, along with the associated predicted costs from the two models using Equation 2.3.

Observed		
Weight(x_i)	Cost(y_i)	Predicted Cost($\hat{y}_i$)
0.001	\$ 4.42	\$ 4.15
0.05	\$ 4.15	\$ 4.15
0.1	\$ 4.15	\$ 4.15
1	\$ 4.15	\$ 4.15
2	\$ 4.23	\$ 4.15
3	\$ 4.41	\$ 4.15
4	\$ 4.57	\$ 4.57
5	\$ 5.13	\$ 5.41
10	\$ 9.39	\$ 9.63
20	\$ 17.65	\$ 18.06
25	\$ 22.67	\$ 22.27
35	\$ 31.51	\$ 30.70
40	\$ 35.10	\$ 34.92
45	\$ 38.70	\$ 39.13
50	\$ 42.93	\$ 43.35

Figure 2.7 displays the observations on a scatter plot and places the regression line over the observations. It shows the resulting regression line plotted through the observed data points. The R-squared value for this model is 0.9975, which represents the fraction of the variance of the

dependent variable that is explained by the model. Note that Figure 2.7 has two regression lines because there are two linear models. If the weight is above 3.5 pounds, the linear model for the data above \$4.15 is used, otherwise the prediction is set to \$4.15.




Figure 2.7: Scatterplot of shipping cost versus weight with a “line of best fit” superimposed over the sample data to show how the two models fit the observed data.

Equation 2.5 provides a model to estimate the shipping cost of an item. In Subsection 2.2.2, we incorporate this model into the main model to track all costs involved with inventory management.

2.2.2 Inventory Management Cost Model Formulation

Here, we formulate a model to study alternative balancing rules to reduce the total cost among the three locations. To compute total cost, the main model uses Equation 2.5 to calculate the associated shipping costs within the system. The following formulation builds the necessary parts of the main model, considering a supply chain inventory system in which demand is fulfilled from Defense Logistics Agency (Defense Logistic Agency (DLA)) warehouses.

This study considers a continuous review inventory policy. Inventory is checked at the end of every day, and an order is placed if the inventory level is less than the reorder level quantity. The MSL of inventory is determined quarterly by the level setting, which is discussed in Section 1.1.

The total inventory management cost for a single item per day comprises holding, backorder, procurement, and lateral transshipment cost terms. In general, the holding cost consists of both

storage fees, which depend on the size and any special handling requirements of the inventory, and the expenses associated with loss, damage, insurances, obsolescence, and the opportunity cost of tying up capital (Silver et al. 1998). In Equation 2.6, the quantity $S_i k_i r_k$ accounts for the expense of storing an item in inventory. This figure is based on the amount of space the inventory occupies in storage and the cost of renting the space. The quantity $S_i p r_p$ accounts for the opportunity cost of tying up capital. This figure is based on the unit price and, like $S_i k_i r_k$, also depends on the amount of inventory. By stocking an item, funding that could have been used for an alternative purpose is sitting on the shelf. There is a chance that the stock may never be used and become obsolete. For the items in this study, the cost of holding inventory in stock for a day is given as

$$C_H = \sum_{i=1}^n S_i k_i r_k + S_i p r_p \quad (2.6)$$

and

$$C_H = \sum_{i=1}^n S_i (k_i r_k + p r_p) \quad (2.7)$$

where

C_H	=	holding costs
n	=	number of locations
S_i	=	average daily stock at location i
k_i	=	cubic foot per item
r_k	=	dollar per cubic foot per day
p	=	unit price
r_p	=	carrying cost percentage per day

The backorder cost arises from stockouts. The system incurs a stockout, or shortage cost, for each unit demanded that it cannot supply from stock (Silver et al. 1998). The price of the stockout does not depend on the unit price, only the price to expedite the inventory. We estimate this expense by assuming the inventory is expedited via two-day shipping. Therefore, the cost of a backorder, based on a two-day emergency shipping cost from Equation 2.5, is given as

$$C_B = \max \left[\sum_{i=1}^n (0.843)w_b B_i + 1.197, \$4.15 \right] \quad (2.8)$$

where

- C_B = backorder costs
- w_b = the weight (*lbs*) of backordered items to be shipped two-day express
- B_i = daily average quantity of backorders at location i

The procurement cost depends on the unit price and ordering cost. The unit price is the price to purchase the item. This includes the shipping paid to the supplier and any additional cost to stock the item (Silver et al. 1998). Within the continental US, shipping is included in the unit price for an item. However, for locations outside the continental US, additional shipping charges may apply in excess of those included in the unit price. The ordering cost is a fixed cost independent of the replenishment size. This cost captures the expense of placing an order (Silver et al. 1998). In this model, the ordering cost is the value of the time spent by the OSPA to queue and approve the PR. The procurement cost, including backorders, is given as

$$C_P = \sum_{i=1}^n Q_i p + A \delta(Q_i) \quad (2.9)$$

where

- C_P = procurement cost
- Q_i = daily average reorder quantity at location i
- p = unit price
- A = fixed cost per order
- $\delta(Q_i) = \begin{cases} 1, & \text{if } Q_i > 0; \\ 0, & \text{if } Q_i \leq 0; \end{cases}$

Many inventory models are based on the economic order quantity (EOQ) model. The purpose of this model is to find the value of the reorder quantity, Q , that minimizes the total relevant costs. Most models built on the EOQ model exclude the first term in Equation 2.9, $Q_i p$, since

this term only depends on demand and not the choice of reorder quantity (Silver et al. 1998). Our inventory model is concerned with the choice of balancing policy rather than the choice of Q . Therefore, the choice of reorder quantity is relevant to our study. The WSS has excess inventory located at locations without demand. Without a balancing policy, this inventory is not used. Implementing a balancing policy will decrease the procurements by using inventory the system would otherwise lose to obsolescence.

The transshipment cost is the expense of shipping inventory from one location to another. The shipping cost for a given weight is the same between any combination of Lemoore, Oceana, and Patuxent River. This is only true if both locations are in the continental US. For other location combinations, a unique shipping model is required to estimate the transshipment cost. From Equation 2.5, the transshipment cost is given as

$$C_T = \max \left[\sum_{i=1}^{n=i} \sum_j^n (0.843)w_i T_{ij} + 1.197, \$4.15 \right], \forall i \neq j \quad (2.10)$$

where

C_T = transshipment costs

T_{ij} = quantity of inventory to be transshipped from location i to j

w_i = weight per item to be transshipped

Note that the relative distance between the two locations does not affect the cost.

Now, the expected cost per period, with transshipment, will be the sum of the holding, backorder, procurement, and lateral transshipment cost:

$$\text{Total Costs} = C_H + C_B + C_P + C_T \quad (2.11)$$

The performance of the system is measured by total cost and the ability to meet demand. There are two ways to measure the ability to meet demand: fill rate and service level. Fill rate describes the proportion of demand met from stock, while service level describes the proportion of all order cycles with no stockout occurrences (Silver et al. 1998).

$$\text{Fill Rate} = \frac{\text{Total Demands} - \text{Total Shortage}}{\text{Total Demands}} \quad (2.12)$$

$$\text{Service Level} = \frac{\text{Number of Periods} - \text{Number of Periods with a Shortage}}{\text{Number of Periods}} \quad (2.13)$$

Generally, both Equations 2.12 and 2.13 can be used to measure a system's ability to meet demand. This study uses Equation 2.12, the fill rate, to evaluate competing balancing policies because this is the preferred performance metric used by the WSS. However, the simulation contains an algorithm from Section 1.2 that sets the MSL and reorder point (ROP) to meet the desired fill rate without direct input on the type of balancing policy in effect.

Our goal is to use Equation 2.11 to compute total costs and fill rates for each of the five NIINs over a two-year period. As a result of the complexities of the analytical modeling and the lack of historical balancing data, we use historical demand and a simulation to estimate the average daily quantity of stock on hand, backorders, reorders, and transshipments. An alternative approach is to find these values by physically altering the actual system. There are two issues with this approach. First, it is costly to experiment with the physical system, and second, each test would take place in real time. Simulations solve these issues by not incurring actual inventory management costs and can repeat the two-year periods using different policies. In Section 2.3, we discuss the simulation model in greater detail.

Table 2.6 shows the variables source and values for the inventory management cost model. Equation 2.7 uses r_k and r_p to calculate holding costs. Equation 2.9 requires A to calculate procurement costs.

Table 2.6: Inventory cost model variables, source, and value.

<i>Variable</i>	<i>Description</i>	<i>Source</i>	<i>Value</i>
r_k	dollar per cubic feet per day	(McNeeley 2013)	0.02114
r_p	carrying cost per day	(Liskow 2013)	0.00069
A	fixed cost per order	(Lasater 2013)	0.72950

2.3 Simulation Structure and Rationale

To evaluate the competing inventory balancing policies, our analysis uses the model summarized in Equation 2.11, formulated in Section 2.2.2, to calculate the total costs and fill rates for an individual policy. These calculations must include the stockout, backorder, reorder, procurement, and transshipped quantities needed. However, historical balancing data is nonexistent, given that a balancing policy has yet to be implemented. As a result of the complexity of the interactions between the random demand and lead times across multiple sites, obtaining these quantities by solving a set of equations may not be possible. This study uses a simulation to generate such quantities, as a simulation can represent the real-world process with fewer restrictive assumptions than those required to solve equations. The simulation provides several two-year estimates of possible “results” within minutes, which would be infeasible through experimentation. Any experiments using the physical system would take years to collect the data necessary to measure the possible impact of competing balancing procedures. Another benefit of a simulation is that it does not interfere with the real system. While there is an inherent risk and cost associated with a real-world experiment, excessive cost within a simulation has no real-world consequences. Furthermore, this study adheres to the specific request by the WSS to provide an analysis prior to physically implementing a balancing policy.

The following analysis uses a discrete event simulation in which the variables change instantaneously and in countable amounts. Time within the simulation is advanced either by events or fixed time intervals (Law and Kelton 2000). This study uses a fixed-increment time-advance mechanism in which the simulation clock advances in increments of exactly Δt , corresponding to one business day. A disadvantage of using the time-advance approach is that it requires the system to check for events after every advance of the simulation clock, regardless of whether any events occur. The inability to skip over periods of inactivity wastes computer resources and lengthens the running time of the simulation. The time-advance approach also forces simulation architects to specify the order in which events occur within the fixed interval (Law and Kelton 2000). For example, the simulation in this study receives outstanding orders prior to the arrival of new demands, and places new orders after all demands for the day have occurred. In reality, a demand could occur before outstanding orders are received and after new orders are placed. See Table 2.7 for an ordered list of events. These disadvantages are negligible in this inventory simulation because events occur in periods of Δt , and the order of events is mandated by the demand and supply planning process within the ERP system.

Most simulations begin with empty systems. For example, an inventory simulation usually begins without any outstanding orders and with the inventory on hand equal to zero, or at the MSL. Since the simulation begins without pending orders, orders are either filled quickly if the system starts with a full inventory, or overly slow if the initial inventory on hand is zero. Information obtained at the beginning of the simulation will influence the quantities of interest, such as the fill rate or the mean holding cost per period, either under- or over-estimating them. This is called an initial bias. Initial bias can be eliminated if data collection begins when the simulation reaches a more representative “steady state.” The period prior to collecting statistical data is called the warm-up period (Law and Kelton 2000).

This study is less concerned with a “steady state” since its interest lies in inventory balancing. If the simulation reaches a steady state, there would be little or no excess to balance. Therefore, our simulation begins with an unbalanced system. This imbalance is achieved by setting the inventory on hand at least equal to the MSL. In addition, the inventory on hand of two of the three locations is set equal to the MSL and a specified number of days of excess inventory is located in the remaining location. The amount of excess for a given number of days is based on the highest average demand data from one of the three locations. A warm-up period of at least 90 days would be necessary to allow the simulation to initially set the reorder period (ROP) and MSL, since level setting is scheduled to occur on a quarterly basis. For a description on the ROP and MSL, and how they fit into the (R, s, S) reorder policy, see the end of Section 1.1.

The simulation includes a 100-day warm-up period, during which no data is collected, followed by a two-year period in which statistical counters gather data. This statistical data include both the cost and quantities of the reorders, backorders, inventory on hand, and transshipments for each simulation day. The number of demands is also collected and used, along with the total backorder quantity to compute the fill rate for the entire two-year simulation period. At the end of the two-year period, the mean of each of these quantities is stored in separate statistical counters after each of the 2,000 replications. Both the warm-up and two-year period occur for every simulation replication. The simulation outputs the mean and standard deviation of these quantities and the cost using the statistical counter for each simulation replication. The total cost for a given inventory transshipment policy is calculated using the mean costs from the main model. The two-year simulation period is long enough to have a sense of the costs associated with each transshipment policy. The impact of the initial excess on inventory management costs becomes negligible in the long term.

The inventory simulation comprises the events listed in Table 2.7. After each advance of the simulation clock, the events occur as they appear in the table, excluding event types three and six, which occurs in quarterly intervals (90 days) and the end of the simulation, respectively. The order of event types four and five is determined by the supply planning process within the ERP system and occur after event types one and two. In reality, event types one and two could overlap; in other words, demand could occur before and after orders are received. The simulation assumes that orders arrive in the morning prior to demand.

Table 2.7: The events that constitute the simulation.

Event Description	Event Type
Arrival of outstanding replenishment order	1
Demand for an item	2
Reset the reorder point (ROP) and maximum stock level (MSL)	3
Redistribute inventory that is in excess to a location that is deficient	4
Place orders for inventory below the ROP up to the MSL	5
End simulation after prescribed duration	6

Demand and lead times do not have deterministic values, but instead are represented by random variables sampled from probability distributions. The daily demand is generated in the following way: One month is selected from the 24 months of the historical demand data. The total demand for the selected month is divided by 25, representing the average number of working days in a month. This value is used as the mean to produce a random Poisson variable, which represents that day's demand. The total demand for the selected month is used in this way for 25 consecutive days within the simulation. On the 26th day, a different month is selected to provide the monthly demand. This process continues until all 24 months are used. To ensure the sequence of the 24 months is random, the order of months is shuffled at the beginning of each replication. Each replication is 600 days long. Figure 2.8 shows the demand per day over the course of a two-year simulation using this process. Each replication uses all 24 months of historical data to generate demand and depicts the number of demands received every day for one simulation replication for NIIN 012030382, a \$1.65 electrical cable. The histogram has two peaks, one on the far left of the plot and another on the far right. The histogram excludes days with zero demand, which accounted for 84 percent of the 600 days in the simulation. Roughly 20 months of the two-year period experience no demand. Three months experience low demand, represented by the first peak, and one month experiences high demand, represented by the second peak.

The lead time for each order is also described by a Poisson distribution. This average lead




Figure 2.8: The following histograms show the number of demands per day for one simulation run (duration = 600 days). The two histograms come from a single item with two obvious data spikes, or modes. There was no demand on 504 of the days. Demands are generated using the monthly total demands from the data. Histogram only includes positive demand, otherwise the histogram would show a spike at zero for 504.

time comes from the data set in Subsection 2.1.4 to determine commonality of demand. The historical mean lead time is used as the mean to generate a random Poisson variable. Unlike the demand, which has a mean that changes each month, the lead time mean is constant throughout the simulation period.

We use common random numbers to ensure each policy is compared using the same demand and lead time inputs for each run. By synchronizing the random demand and lead time inputs across the competing policies, it becomes easier to distinguish the differences between the policies. Without using the same demand and lead times, the differences in competing policies may be due to the variation in the inputs and not the policies themselves.

The simulation uses a method discussed in Ross (2006) to generate a random Poisson variable. The mean from the historical data is represented by the Greek letter λ . The simulation generates a sequence of uniform independent and identically distributed $(0, 1)$ variables, U , until there is one random variable such that $U_i < e^{-\lambda}$. The value i , such that $U_i < e^{-\lambda}$, is the random Poisson variable. For example, if the mean demand is 4 and the sixth random U generated is less than e^{-4} , then the random variable representing demand for that day is 6 (Ross 2006). Table 2.8 provides the attributes of the Poisson distribution. In the upcoming event discussion, it is important to

remember that only a single item is simulated at a time.

Table 2.8: The Poisson distribution. (After Wackerly, Mendenhall, and Scheaffer 2008)

Poisson	Poisson(λ)	
Application	Models the probability distribution of the occurrence of rare events in space, time, or volume, where λ is the average value of X .	
Mass	$p(x) = \frac{\lambda^x}{x!} e^{-\lambda}$,	$x = 0, 1, 2, \dots, \lambda > 0$
Cumulative	$F(x) = e^{-\lambda} \sum_{i=0}^x \frac{\lambda^i}{i!}$,	$x = 0, 1, 2, \dots, \lambda > 0$
Mean	λ	
Variance	λ	

In this study, we model an inventory system. A system is defined as a group of objects whose actions and interactions accomplish a specific task. State variables describe the system at a given time (Law and Kelton 2000). These state variables are necessary to produce the simulation results for our analysis. Examples of state variables within the system include the inventory on hand at each location, the collection of outstanding orders and their arrival dates, and all other variables necessary to describe the inventory system. The events listed in Table 2.7 impact the state variables. For example, in the first event, an arrival of an outstanding order changes both the inventory on hand and the outstanding order state variables.

2.3.1 Fixed Events

In the following discussion, we describe the logic for each event and how it impacts the system's state variables. In Subsection 2.3.1, we focus on the events that do not change, regardless of the balancing policy. In Subsection 2.3.2, we depict the redistribution event, but only the WSS proposed policy is discussed in detail.

The flowchart in Figure 2.9 illustrates the replenishment order arrival event, including the update to state variables resulting from this process. The simulation performs this event first at the beginning of each simulation day. The simulation uses a priority queue to return the earliest pending orders by implementing a heap data structure. The benefit of a heap data structure is that it returns the minimum value in big O notation, $O(1)$. However, it takes $O(\log n)$ to insert and delete orders in the queue. For each simulated business day, as shown in Figure 2.9, the replenishment arrival event checks if the minimum delivery date equals the current simulation date. If these dates are not equal, then the function is complete. However, if they are equal, the inventory on hand is incremented by the amount of the order, and the order is removed from the

heap. It is also possible for multiple orders to arrive on the same day. Therefore, after removing the received order from the heap, the function checks again if the minimum delivery date matches the current simulation date. This continues until the minimum delivery date no longer matches the current date. In the simulation, backorders are allowed and are represented by a negative inventory on hand. When an order is received, it fills the backorder before fulfilling the inventory shortage, if it exists.




Figure 2.9: The following flowchart diagrams the order arrival event within the simulation.

A flowchart for the demand event is given in Figure 2.10, and shows the resulting changes to the simulation's state variables and statistical counters.




Figure 2.10: The following flowchart shows the demand function within the simulation.

First, the simulation uses the common random demand input for the given day and replication. This ensures that each balancing policy faces the same random demand. The simulation records the number of stockout occurrences by comparing the demand requirement to the balance on hand. Any requirements in excess of the quantity on hand are recorded in a statistical counter as an out-of-stock occurrence. The inventory on hand state variable is then decremented by the amount of the demand input. The inventory on hand is allowed to be negative to represent backorders. If the inventory on hand is negative at the beginning of the event, then all demands will result in a stockout occurrence.

The flowchart in Figure 2.11 shows the reset reorder point and maximum stock level events, which both take place every 90 days.




Figure 2.11: The following flowchart diagrams resetting the reorder point (ROP) and maximum stock level (MSL).

Silver, Naseraldin, and Bischak's (2008) algorithm determines the ROP and MSL for a (R, s, S) policy to meet a desired fill rate. This algorithm accounts for the delay between when the inventory on hand falls below the reorder point and when the next reorder review occurs. This delay requires that the s and S in a (R, s, S) policy be greater than that of a continuous review policy to meet the same desired fill rate. Within the simulation, their algorithm is relatively robust in meeting the desired fill rate, as long as the inventory does not radically vary from month to month. Unfortunately, the algorithm has trouble maintaining a given fill rate when demand is lumpy, for the reasons discussed previously in Section 2.1. Although the algorithm may perform poorly for some of the demand patterns, the purpose of this study is to compare balancing policies. As long as the algorithm is consistent across all policies, we can use it to

compare the balancing policies.

The decision to reorder stock, shown in Figure 2.12, is examined at daily review intervals. If the inventory position (on hand plus on order) is at or below the ROP, s , then the amount required to bring the inventory level up to the MSL, S , is ordered. The review interval, R , depends on the type of inventory and the location (daily for the Trident Refit Facilities (TRF), and twice a week for NASs and Fleet Logistic Centers (FLCs)). The review period is every business day. The order is assigned a delivery date and added to the heap.




Figure 2.12: The following flowchart diagrams the placing an order event within the simulation.

2.3.2 WSS Proposed Balancing Policy

In the previous flowcharts, the events occur independently at each of the locations. For example, an item received at Oceana does not affect Lemoore. The redistribution event is different because excess at one location is transshipped to another site. This interaction causes state variables to be updated across multiple locations simultaneously. To sort and compare locations, they are stored in collections called bins. Figure 2.13 illustrates the sorting process for each location.

The redistribution event is more complex, because the sites interact with each other, unlike previous events. As a result of this complexity, the study breaks down the redistribution event into multiple stages. The first stage separates the locations into three separate bins using the

Compare method: equal bin, excess bin, and shortage bin. For example, if the MSL minus the inventory on hand is less than zero, the location is sorted into the excess bin, while if it is greater than zero, it is sorted into the shortage bin. The shortage bin sorts the locations by the largest shortage first, while the excess bin sorts by the largest excess first. Figure 2.13 illustrates this initial stage.




Figure 2.13: The following flowchart diagrams the first stage in the redistribution event.

The redistribution event ends if there is no excess, since there is nothing to balance. Figure 2.14 provides a flowchart of this stage. If there is an excess, the flowchart selects the (next) location with the next largest excess. The word “next” is in parentheses, indicating that on the first pass through the flowchart, the location with the largest excess inventory is chosen first. With each additional pass through stage two, the location with the next largest excess is chosen. Since the excess bin is sorted from largest excess to smallest, this means the function starts with the location on the top of the bin, and with each subsequent pass through stage two, selects the next lower location in the bin.

As in the second stage, the redistribution event can end in the fourth stage if there are no more shortages to balance, as there is no longer a need to balance. The flowchart shown in Figure 2.15 is similar in structure to Figure 2.14, but rather than iterating through the excess bin, the third stage iterates through the shortage bin. Since the shortage bin is organized from greatest shortage to least, each pass through stage three chooses the next greatest shortage.




Figure 2.14: The following flowchart diagrams the second stage in the redistribution event.

While the second stage in the redistribution event selects an excess location, illustrated in Figure 2.14, the third stage selects a shortage location, as shown in Figure 2.15. The fourth stage compares these two selections. If the excess and shortage sites are not in the same balancing region (i.e., their unit price is below the threshold to balance between the two locations), the function returns to the third stage and selects the next highest shortage. It then returns to the fourth stage, assuming another shortage location exists, and compares the two new selections. This continues until the function has cycled through all shortage locations in the bin and as long as there is excess inventory to distribute at the selected excess location. The function then returns to the second stage and selects the next excess location in the bin. If the excess bin is empty or if all excess locations have been tried, the redistribution event ends.

The fourth stage is illustrated in Figure 2.16. If there is a permissible excess and shortage combination, then the state variables are incremented and decremented as per Figure 2.16. A function determines if the inventory on hand in the shortage locations is below its MSL. A shortage may remain at the shortage location even after all inventory has been received from the excess location. If this occurs, the excess region is removed from the excess bin, and the function returns to the second stage to select the next excess location, if one exists. In contrast, if




Figure 2.15: The following flowchart diagrams the third stage in the redistribution event.

the shortage location no longer has a shortage after receiving inventory, its inventory on hand is set equal to its MSL and it is removed from the shortage bin. Any excess inventory above the MSL of the shortage location is returned to the excess location. The function then returns to stage three to select the next shortage location.

The redistribution event is unique in that it changes based on the implemented inventory balancing policy. The previously discussed stages describe the inventory balancing policy proposed by the WSS. The goal of this study is to evaluate different balancing policies, requiring the logic of the redistribution event to change based on the policy in use. For example, an alternative balancing policy may require the excess and shortage bins to be sorted by geographic location rather than largest excess and shortage. The three initial policies this study tests are one that does not balance, one that uses the thresholds and regions proposed by the WSS, and one that balances without thresholds.

2.3.3 Statistical Counters

The statistical counters for the Q_i , S_i , T_{ij} , and B_i quantities are recorded each day throughout a simulation run. The average of each of these statistical counters is calculated at the end of each




Figure 2.16: The following flowchart diagrams the fourth stage in the redistribution event.

replication and stored in a statistical counter. At the end of all the simulation runs, the average and standard deviation of these quantities are the output required to run the total cost model. At the end of each simulation run, the fill rate is calculated from the demand statistical counter and the number of stockouts statistical counter. The fill rate is also recorded in a statistical counter at the end of each simulation run. The average and standard deviation of the fill rate over all the simulation runs is another output of the simulation.

2.3.4 Single Location, Single Day Flowchart

Table 2.12 depicts a single day within the simulation. The balance algorithm changes depending on the policy and is not shown in detail.




Figure 2.17: The following flowchart diagrams a single day within the simulation.

CHAPTER 3:

Analysis

The study uses a simulation to evaluate three different transshipment policies to determine which delivers the best results in terms of fill rates and costs. The first policy, *no lateral shipments (NLS)*, does not allow for materiel to be moved among stocking locations and acts as a baseline for comparison with the other policies. The second policy, *proposed*, allows for transshipments among stocking locations if they meet the unit price thresholds proposed by Weapon Systems Support (WSS). The third policy, *no threshold*, allows transshipments regardless of unit price if they meet quarterly volume requirements. Our analysis will only require a simulation to be run for two of the three policies for each item. This is because two of the policies behave in the same manner and produce the same simulation results; for example, the *proposed* policy acts in the same manner as the *NLS* policy and the *no threshold* policy. Since unit price is either below or above the balancing threshold and there is demand at only two of the three locations in all but one case, there is no unique instance in which all three policies behave differently. Therefore, we only compare between balancing and non-balancing policies.

Two of the policies are always exactly the same when demand occurs at only two of the three locations. For items below the threshold, under the proposed policy, the *NLS* policy is the same as the *proposed* policy since in either case, balancing would not occur. For items above the unit price threshold, the *proposed* policy is the same as the *no threshold* policy since in either case, balancing would occur. In the single case in which there is demand at all three locations, the unit price is above the unit price threshold to allow balancing within the continental United States. In this case, the demanded item is allowed to be balanced within the continental United States, in which case the *proposed* policy is again the same as the *no threshold* policy. Therefore, we are really comparing between balancing and non-balancing policies since two of three policies always behave in exactly the same manner.

To determine the impact of the three policies, we consider five items held at three different locations. For each item, we evaluate a balancing policy and compare it with one in which there is no balancing. To do this, we track and record the resulting data on fill rates, holdings, backorders, transshipments, procurements, and total relevant costs. Due to the complexities involved in moving materiel across locations, a simulation is used instead of an analytical model. The simulation includes 2,000 replications to calculate mean reorder cycle quantities. In this

study, the simulation spans a two-year period and begins with an initial excess of 30 days. This excess is based on the average historical demand of the location with the highest demand.

We begin the simulation with excess that already exists within the system, because without excess, the impact of inventory balancing would be negligible and not representative of the \$794 million dollars of excess currently in the system. Therefore, our simulation begins with an unbalanced system. We create this imbalance by first setting the on-hand inventory equal to the maximum stock level (MSL) for all three locations and then setting one of the three location's on-hand inventory levels equal to the MSL plus a specified number of days of excess. To assess the balancing rules, we vary the location of the excess materiel. We first consider the impact of balancing when the initial excess resides at the high-demand location, then the moderate-demand location, and finally the low-demand location.

The locations with low demand in fact have zero demand for four of the five examined items over the duration of the simulation. One item is demanded at all three locations. Excess at a location without demand is representative of the data and its inclusion in the analysis is important. One explanation for this excess is the drop in demand that results from homeport changes or the decommissioning of a platform at a particular site. Although demand may no longer exist at one site, other sites may still require the materiel under consideration.

3.1 Initial Results

Table 3.1 shows the differences in system-wide fill rates as well as the percent changes and dollar differences in inventory management costs resulting from the implementation of balancing for each of the five examined items when the excess resides at the high-, moderate-, and low-demand locations. A percent change measures differences in values and usage percentages and conveys the magnitude of the changes resulting from the balancing policies. This study presents differences in the fill rates and total costs of all three locations before and after the implementation of balancing. An individual location's fill rate and cost percent may move in the opposite direction of the overall percent change for all three locations. A positive percentage denotes an improvement in value, whereas a negative percentage denotes a degradation in value.

For all five examined items, the balancing policy results in lower relevant costs when the excess resides at the location with zero demand for the item. This is because, without balancing, the excess incurs a constant holding cost at the location with zero demand and the locations with demand must procure the same materiel. However, when the excess resides in the high-demand

Table 3.1: Percent changes represent differences in the fill rates and costs between balancing and not balancing. **A positive percentage indicates improvement, while a negative percentage indicates a decrease in value after balancing implementation.**

Location of Excess	% Improvement		Daily
Item Description	Fill Rate	Total Cost (%)	Total Cost (\$)
Location of Initial Excess			
High Demand Location			
VALVE, REGULATING	< -0.01 %	< -0.01 %	0.45
WIRE, ROPE	-0.09 %	0.18 %	0.78
SEAL, BOSS	-0.58 %	0.45 %	0.32
HOOK POINT	0.00 %	0.16 %	2.42
CONE AND ROLLERS	< -0.01 %	-0.03 %	-0.05
Location of Initial Excess			
Medium Demand Location			
VALVE, REGULATING	1.96 %	16.48 %	25,839.92
WIRE, ROPE	-3.11 %	1.32 %	5.71
SEAL, BOSS	-0.14 %	19.89 %	17.66
HOOK POINT	< -0.01 %	-0.01 %	-0.09
CONE AND ROLLERS	-0.01 %	3.80 %	6.18
Location of Initial Excess			
Low or No Demand Location			
VALVE, REGULATING	1.96 %	16.48 %	25,839.91
WIRE, ROPE	-0.81 %	4.23 %	18.90
SEAL, BOSS	-0.11 %	20.940 %	18.86
HOOK POINT	0.01 %	1.349 %	20.13
CONE AND ROLLERS	-0.01 %	3.94 %	6.43

location, the impact is small and, in some cases, negative. When the excess resides in the moderate-demand location, the impact is usually greater than that at the high-demand location but lower than that at the low-demand location.

In addition to inventory management costs, we also compare differences in fill rates before and after the implementation of balancing. In our simulations, the system-wide fill rate rarely improves. Although fill rates may increase at individual locations, the system's overall fill rate, in general, actually decreases. We examine this phenomenon in greater detail in Subsection 3.1.1.

Upon review, the results in Table 3.1 raise some interesting questions: When the excess resides in the moderate-demand location, why is its percentage change value not between the values found at the high- and low-demand locations? Why does the implementation of balancing considerably improve the inventory management costs for the VALVE, REGULATING and SEAL, BOSS items when excess resides in the low- and moderate-demand locations? Finally, the

most counterintuitive result raises the biggest question: Since the anticipated benefit of inventory balancing is an increase in fill rates, why does the balancing result in a decrease in fill rates in 11 of the 15 examined scenarios? Some of the factors driving these counterintuitive results include the lumpy nature of demand, low demand, and frequency of level setting, and will be discussed in greater detail in the following sections.

First, it is important to understand how the values in Table 3.1 are calculated; more importantly, however, it must be recognized that each percent change is specific to the nature of excess at a particular location before and after balancing. The percent changes in Table 3.1 use Equation 3.1 and Equation 3.2. Equation 3.1 calculates differences in fill rates, which are shown in the first column in Table 3.1.

$$\text{after value} - \text{before value} = \text{difference} \quad (3.1)$$

If the initial excess resides in the moderate-demand location for the WIRE, ROPE item, the overall fill rate across all three sites is 87.54 percent before the implementation of balancing and 84.43 percent after balancing implementation, using Equation 3.1 to calculate the difference.

$$\text{after value} - \text{before value} = 84.43 - 87.54 = -3.11\%$$

To show a positive change with a reduction in cost, the values in Equation 3.1 are rearranged to show this sign convention. Equation 3.2 calculates the percentage change in cost.

$$\frac{\text{before value} - \text{after value}}{\text{before value}} \times 100 = \% \text{ change} \quad (3.2)$$

Again, if the excess resides in the moderate-demand location for the WIRE, ROPE item, its total procurement cost across all three sites is \$433.75 before the implementation of balancing and \$428.03 thereafter, using Equation 3.2 to calculate the percent change.

$$\frac{\text{before value} - \text{after value}}{\text{before value}} \times 100 = \frac{\$433.75 - \$428.03}{\$433.75} \times 100 = 1.32\%$$

The fourth column in Table 3.1 shows this improvement in procurement costs resulting from the

implementation of balancing for the WIRE, ROPE item.

While Table 3.1 shows the overall inventory management costs from balancing for each of the five items, Figure 3.1 breaks down the total costs for the BOSS, SEAL item with and without balancing. Notice that the procurement cost is the primary driver of whether balancing improves inventory management costs. Also notice that inventory management costs are relatively constant under balancing, regardless of where the initial excess resides. However, inventory costs increase when the initial excess moves from a location with high demand, Patuxent River, to a location with no demand, Lemoore, under the NLS policy.




Figure 3.1: The bar chart shows the total costs per day for the SEAL, BOSS item with and without balancing, where the simulation begins with 30 days of excess in the indicated location. The colors within each bar indicate the portion of holding, backorder, procurement, and transshipment costs that contribute to each total cost.

As seen in Figure 3.1, holding costs are the second-largest contributors to inventory management costs. However, there is a relatively small difference in holding costs under balancing and not balancing in relation to procurement and backorder costs, except when the excess resides in a location with little to no demand. In general, holding costs only impact the decision to balance

when the excess resides in the low-demand location. However, holding costs impact the decision to balance for the SEAL, BOSS item when excess resides in both the moderate-demand location, Oceana, and the zero-demand location, Lemoore. This is because Oceana’s demand is relatively low compared to the amount of excess. Finally, notice that transshipment costs make up the smallest portion of inventory management costs. Figure 3.1 is representative of the makeup of inventory management costs in general.

Both balancing policies evaluated in this study are preventive. Preventive transshipment is an attempt to move materiel from a location that is unlikely to experience a stock-out to a location that is more likely to face this situation. We anticipate that the implementation of a preventive transshipment policy will produce benefits. First, fill rates will increase because the materiel is located where demand is most likely to occur. Second, materiel storage costs will decrease because materiel is moved to where it will be used instead of remaining at a location where the on-hand inventory exceeds its MSL. Third, procurement costs will decrease because the excess at other locations will offset the need to procure materiel. However, some of the simulation’s results contradict the anticipated benefits of inventory balancing. This is because these anticipated benefits assume that the materiel will be moved to the location where demand for it is most likely. However, the business rules governing transshipment do not always move materiel to the location that is most likely to experience demand for it. For example, fill rates decrease for four of the five examined items when the excess resides in the high-demand location. Table 3.2 lists the anticipated and observed costs and benefits.

Table 3.2: Anticipated versus observed costs and benefits from transshipment.

<i>Anticipated</i>		<i>Observed</i>	
Cost		Cost	
Transshipment costs from inventory balancing		Transshipment costs are a small percentage of total costs Highest observed percentage is 0.22%	
Benefits		Benefits	
Improved ability to meet demand		Observed improvement to meet demand in 4 of 15 scenarios	
Savings from reduced holding costs		Observed decrease in holding costs for all 15 scenarios	
Savings by preventing procurements		Observed savings in procurement costs in 14 of 15 scenarios	
Savings by preventing loss due to obsolescence			

We specifically focus on these departures from the anticipated benefits of transshipment, including cases in which fill rates decrease and backorder costs increase, in the following sections.

3.1.1 Departure from Anticipated Benefits: Fill Rate Decreases

A purported benefit of inventory balancing is an improved ability to meet demand, measured by a higher increase in fill rate. However, in Table 3.1, three of the five items show some decrease in fill rate when balancing is enabled, regardless of where the initial excess resides. It is expected that such redistribution would decrease the number of stockout occurrences and would, therefore, increase the fill rates. There are two reasons why this decrease in fill rate occurs: First, due to a location's low demand, its MSL is set to zero for a large portion of the simulation time. The excess that is transferred to this location fulfills a demand, preventing a procurement. However, since the MSL is zero, it does not prevent a stockout occurrence. Therefore, the location in which this initial excess resided has less materiel to prevent its own shortages in the future, which can decrease the system's overall fill rate. In Subsection 3.1.2, we investigate why this occurs through a comparative analysis of preventive and emergency transshipment policies. Second, quarterly level setting is too infrequent to handle extreme fluctuations in demand for some items. Level setting adjusts the MSL based on past demand, but these demands do not always follow consistent patterns, meaning that zero or very little demand occurs for many months followed by a month of very high demand. If the past demand is high, it sets a high MSL. However, a period of high demand is often followed by a period of no demand. Therefore, level setting actually ends up setting a high MSL when future demand is low and setting a low MSL when demand is high. For example, when the MSL is low and demand is high, excess items are shipped to the other locations. Since this is a period of high demand, there are no items left to meet it and, therefore, stockout occurrences increase. This contrast between future demand and the set MSL leads to a decrease in fill rate when balancing is enabled.

3.1.2 How Transferred Excess Fails to Prevent Shortages

As discussed in 1.2, inventory balancing is categorized by two types of transshipment: preventive and emergency. Currently, inventory balancing within Navy Enterprise Resource Planning (ERP) is a form of preventive transshipment, as are the policies examined in this study, but the observed differences between emergency and preventive transshipment help identify why fill rates decrease in our results when one location's MSL is zero for the majority of the simulation.

Emergency transshipments expedite materiel only after demand occurs, whereas preventive policies call for materiel to be shipped at any time when the on-hand inventory is lower than the MSL. With emergency transshipment, stock-out is always prevented due to the way in which stock-out occurrences are defined. This is because an emergency transshipment policy

considers the shipped materiel from another location as meeting demand rather than as a stock-out occurrence. In contrast, a preventive transshipment policy defines a stock-out occurrence as any time when materiel is not on hand at the time of demand, regardless of whether excess is used to fulfill the demand. When the on-hand inventory is zero, these two policies behave in exactly the same manner. Even when demand occurs at a location with an on-hand inventory of zero, both emergency and preventive transshipment policies would call for any excess to be shipped to the location with the demand. However, the real difference between the two policies is that emergency transshipment does not consider this situation to be a stock-out occurrence since the excess was used to meet demand, whereas preventive transshipment does.

An on-hand inventory of zero is rare when the locations have an MSL above zero. However, as long as the MSL is zero and the site has no excess, the on-hand inventory will be zero. In this case, the transshipped materiel may still lower holding and procurement costs, but it does not increase fill rates. The Navy will still ship excess to a site after demand occurs as long as the on-hand inventory is below the MSL, including an MSL equal to zero. The CONE AND ROLLER and WIRE, ROPE items illustrate this behavior.

WIRE, ROPE (NIIN 011669450)

This is the third most-expensive item, with a unit cost of \$81.23. It falls within the second-tier threshold in the WSS proposed business rules and is eligible for balancing within a regional area (NE, SE, NW, SW, Trident, Japan less Guam). Regardless of where the initial excess resides, the overall fill rate for this item decreases. Over the two-year period, the WIRE, ROPE item at Oceana, the moderate-demand location, has a mean demand of 76, with a standard deviation of nine. For such low levels of demand, level setting adjusts the MSL to zero at Oceana for the majority of the simulation time.

We previously examined the overall fill rate among all three locations together. Table 3.3 breaks down the fill rates for each of the individual locations. The first policy, *NLS*, does not allow for the movement of materiel across stocking locations. Again, due to its unit price, transshipments under the WSS proposed balancing policy would be allowed between Oceana and Patuxent River but not between East and West Coast locations. However, under the third policy, *no threshold*, transshipments are allowed among all stocking locations. When excess resides at the high-demand location, Patuxent River, the fill rate declines by 0.18 percent, but it improves at the moderate-demand location, Oceana, by 1.54 percent. The net impact is an overall decrease in the fill rate because of the decrease at Patuxent River, which has greater demand. Due to its

Table 3.3: Fill rate for each location for NIIN 011669450 after running the simulation 2,000 times with an initial 30 days of excess at the high-demand location.

<i>Balancing</i>	<i>Location</i>	<i>Initial Excess</i>	<i>Fill Rate (%)</i>	
			μ	σ
<i>NLS</i>	Patuxent River	<i>Yes</i>	86.52%	10.99%
	Oceana	<i>No</i>	55.65%	14.55%
	Lemoore	<i>No</i>	100.00%	0.00%
<i>No threshold</i>	Patuxent River	<i>Yes</i>	86.34%	11.17%
	Oceana	<i>No</i>	57.09%	15.89%
	Lemoore	<i>No</i>	100.00%	0.00%

greater demand, Patuxent River (*high*) has a greater influence over the fill rate. Among all three locations, the fill rate decreases in 60.75 percent of the 2,000 replications. In 4.4 percent of the replications, there is no difference in the fill rate, while in 34.85 percent of the replications, the fill rate improves. Performing a matched pair *t* test confirms that the differences in the fill rates are statistically significant.

Inventory balancing within Navy ERP first meets demand by sending excess to the location with deficiencies prior to procuring the materiel. Figure 3.2 illustrates this redistribution when the receiving location's MSL is zero. Again, balancing only improves fill rates when materiel




Figure 3.2: Inventory balancing within Navy ERP screens locations for materiel above the MSL to be redistributed to locations with inventory below the MSL. After redistributing materiel between locations, the Navy ERP places orders to fulfill the remaining inventory deficits.

are moved from a location that is unlikely to experience stock-outs to a location that is more

likely to face this situation. Since Oceana (*moderate demand*) does not have an allowance and is characterized by low and infrequent demand, level setting adjusts its MSL to zero for the majority of the simulation. In this situation, Oceana (*moderate demand*) experiences a stock-out before materiel can be transferred from Patuxent River (*high demand*). However, the excess materiel still prevents procurement from taking place at Oceana (*moderate demand*). This is also the case when the initial excess resides at the moderate- and zero-demand locations.

Reduction in the Impact of an MSL of Zero

Since Oceana's MSL is set to zero for the majority of the simulation, excess materiel do not improve fill rates under balancing, as previously discussed. Rerunning the simulation with an allowance increase from zero to two at Oceana (*moderate demand*) results in a 0.01 percent increase in fill rates while still reducing total costs. Performing a similar analysis on the CONE AND ROLLER item, we find similar behavior to the WIRE, ROPE item when the allowance is increased.

3.1.3 How Poor Level Setting Decreases Fill Rates under Balancing

As previously discussed in Section 3.1, preventive transshipments attempt to reduce the risk of possible future stock-out occurrences by moving materiel to where they will most likely be needed in the future. The difficulty is in accurately predicting the location where future demand will occur. Navy ERP uses the on-hand inventory's relationship to the MSL to determine where materiel should be moved. This approach depends on an accurate MSL to function correctly. Erratic demand can lead to a high MSL when future demand will be low. In this case, inventory balancing will be unable to move materiel to other locations where they will more likely prevent stock-outs. Likewise, if level setting adjusts to a low MSL when future demand is high, inventory balancing will transfer materiel away from locations that will experience high demand soon to locations that are less likely to experience stock-outs. In both cases, inventory balancing will decrease fill rates compared to not balancing. The VALVE, REGULATING item illustrates this behavior.

VALVE, REGULATING (NIIN 012030382)

This item has the highest dollar volume of the five items under consideration. It falls within the second-highest threshold in the WSS proposed business rules and is eligible for balancing within the continental US and between Guam and Japan. When excess resides in the high-demand location, this item's overall fill rate decreases. Oceana's demand is erratic, with a mean daily demand of 82.5 over a two-year period and a standard deviation of 635.1. Lemoore's daily

demand is relatively stable, with a mean demand of 0.004 and a standard deviation of 1.5. Both moderate- and high-demand locations have allowances of two. Figure 3.3 shows the original historical data for the VALVE, REGULATING item, except during September 2012. This is because the demand in September 2012 was 59,365 and, if included in the chart, would make the other monthly demands indistinguishable from zero.




Figure 3.3: Bar chart of monthly demands at Oceana from October 2010 to August 2012 for the VALVE, REGULATING item. The demand for September 2012, of 59,365, is not shown because it would make the bar chart for the other months indistinguishable due to scaling. (After Liskow 2013)

According to Table 3.4, the first policy, *NLS*, does not allow for materiel to be moved across stocking locations. However, this item meets the requirements of the second policy, *proposed*, which allows for transshipments from all stocking locations to all other stocking locations for items with unit prices of \$75.00 and above. As shown in Table 3.4, the fill rate decreases slightly at the excess location of Oceana (*high demand*), but it improves at Lemoore (*moderate demand*). The net impact is an overall decrease in the fill rate. Since Oceana (*high demand*) has a greater demand, it has a greater impact on the fill rate than Lemoore (*moderate demand*). In 45.3 percent of the 2,000 replications, there are no differences between the fill rates. In 3.3 percent of the replications, balancing improves the fill rate, while in 51.4 percent of replications, the fill rate decreases.

Although the differences in fill rates before and after balancing seem minor, performing a matched pair *t* test shows that there is sufficient evidence that these differences are statistically significant. The erratic demand at Oceana (*high demand*) leads to an MSL that is low when future demand

Table 3.4: Fill rate for NIIN 012030382 at each location after running the simulation 2,000 times with an initial 30 days of excess at the high-demand location.

<i>Balancing</i>	<i>Location</i>	<i>Initial Excess</i>	<i>Fill Rate (%)</i>	
			μ	σ
<i>NLS</i>	Patuxent River	<i>Yes</i>	100.00%	0.00%
	Oceana	<i>No</i>	28.73%	20.21%
	Lemoore	<i>No</i>	88.92%	20.27%
<i>Proposed</i>	Patuxent River	<i>Yes</i>	100.00%	0.00%
	Oceana	<i>No</i>	28.72%	20.20%
	Lemoore	<i>No</i>	96.91%	11.45%

is high. In this situation, Oceana (*high demand*) ships excess materiel to Lemoore (*moderate demand*) even though it is more likely to face shortages due to its MSL setting. Therefore, the fact that Oceana's on-hand inventory is below or above the MSL is a poor indicator of which location is more likely to face shortages. This is in part because infrequent level setting is unable to keep pace with the erratic demand at Oceana (*high demand*). Lemoore (*moderate demand*) maintains a high fill rate prior to implementing balancing. This illustrates why the fill rate decreases when balancing is implemented because materiel are not transferred to where the next shortage is more likely to occur.

Increase the Frequency of Level Setting

Since the MSL is a poor indicator of where the next shortage will occur due to infrequent level setting at Oceana (*high demand*), we will rerun the simulation, increasing the frequency of level setting from every 90 days to every two days. This results in a 44 percent increase in the fill rate under balancing. Table 3.5 shows the fill rate results when initial excess resides in the high-demand location and the frequency in level setting is increased from every 90 days to every two days.

Under the *NLS* policy, the fill rate improves with the higher level setting frequency because the MSL more often reflects demand and locations are therefore less likely to experience stock-outs. Notice that the fill rate for the moderate-demand location, Lemoore, changes by less than one percent compared to the original level setting frequency, per Table 3.4 and Table 3.5. Increasing the level setting frequency improves the system's fill rate. This is because the MSL's relationship with the on-hand inventory level becomes a good indicator of where the next shortage will occur.

Table 3.5: Fill rate for NIIN 012030382 for each location after running the simulation 2,000 times with an initial 30 days of excess at the high-demand location (increased level setting).

<i>Balancing</i>	<i>Location</i>	<i>Initial Excess</i>	<i>Fill Rate (%)</i>	
			μ	σ
<i>NLS</i>	Patuxent River	<i>Yes</i>	100.00%	0.00%
	Oceana	<i>No</i>	35.25%	41.17%
	Lemoore	<i>No</i>	89.93%	20.21%
<i>Proposed</i>	Patuxent River	<i>Yes</i>	100.00%	0.00%
	Oceana	<i>No</i>	50.83%	34.54%
	Lemoore	<i>No</i>	97.17%	11.19%

Performing a similar analysis on the SEAL, BOSS item shows that it behaves like the VALVE, REGULATING item when the level setting frequency is increased.

3.1.4 Departure from Anticipated Benefits: Increase in backorder Costs

Similar to our expectation that fill rates would increase under balancing, we assume that backorder costs will decrease. The reasons that backorder costs actually increase should be analogous to the discussion in Section 3.1.1 on why fill rates decrease. However, the backorder costs for the SEAL, BOSS item actually improve as fill rates decrease. In this study, backorder costs are based on the cost of expediting materiel. No matter how little the shipment weighs, there is a minimum cost of shipping materiel. In the case of this item, the costs of shipping one or 25 pieces are the same. The shipping of excess materiel to Oceana prevents stock-outs at this location. Oceana's stock-out occurrences are less frequent than at the high-demand location, Patuxent River, resulting in greater expediting costs. Although Patuxent River experiences more stock-out occurrences when it ships its excess materiel to Oceana, stock-outs at Patuxent River tend to occur on the same day. This allows expedited stock-out orders to be combined, resulting in lower costs. Therefore, although fill rates decrease with balancing, backorder costs decrease for this item.

3.1.5 Departure from Anticipated Benefits: Increase in Procurement Costs

Only one case results in the increase of procurement costs under balancing (i.e., NIIN 014617380). This increase is slight, at only 0.072 percent. It occurs because Oceana is more likely to reach its reorder point (ROP) when it transfers its excess to the other locations. This explains the slight increase in procurement costs under balancing, but it is not considered a significant issue.

3.1.6 Minimal Impact of Transshipment Costs

Table 3.6 shows that transshipment represents less than one percent of total cost in all 15 scenarios. In all cases, transshipment costs make up a larger proportion of total cost when the excess resides at the low-demand location. The performance of the various balancing policies is driven primarily by inventory holdings, procurement, and fill rates. The prevention of churn, although important, seems like an ancillary consideration with regard to whether to balance inventory.

Table 3.6: Percent of total costs due to transshipments.

<i>ITEM/Excess Location</i>	<i>% of Total Cost</i>	<i>ITEM/Excess Location</i>	<i>% of Total Cost</i>
VALVE, REGULATING		WIRE, ROPE	
Patuxent River	0.018 %	Patuxent River	0.033 %
Oceana	0.002 %	Oceana	0.072 %
Lemoore	0.016 %	Lemoore	0.078 %
SEAL, BOSS		HOOK POINT	
Patuxent River	0.126 %	Patuxent River	0.004 %
Oceana	0.146 %	Oceana	0.003 %
Lemoore	0.219 %	Lemoore	0.002 %
CONE AND ROLLERS			
Patuxent River	0.017 %		
Oceana	0.170 %		
Lemoore	0.172 %		

3.2 Sensitivity Analysis

In this section, we perform a sensitivity analysis in which we determine how different parameter values impact fill rates and inventory management costs. The first parameter we focus on is the amount of excess. We vary the initial excess quantity at Patuxent River for the WIRE, ROPE item to see if our original insights hold. We then vary the mix of locations. Thus far in our analysis, we have only considered naval air stations (NASs). In this section, we replace NASs Oceana with Naval Submarine Base Kings Bay. With this change of location, we will need to reselect the five items under consideration to ensure that they are commonly demanded by at least two of the locations. Again, we will use dollar volume to select these items. We perform the same analysis as previously applied to the five original items. The relatively high allowance for Kings Bay compared to NASs focuses our sensitivity analysis on allowancing.

3.2.1 How the Amount of Initial Excess Impacts Fill Rates and Costs

Figure 3.4 shows both the fill rates and inventory costs for the WIRE, ROPE item before and after the implementation of balancing when the initial excess resides at the high-demand location. The unit price for this item is \$81.23 and it is eligible for balancing within a regional area under the proposed WSS policy. Since the moderate-demand location, Oceana, has an MSL equal to zero for the majority of the simulation, excess shipped from the high-demand location, Patuxent River, does not improve Oceana's fill rate.

For various amounts of initial excess, Subfigure 3.4a shows that the fill rates before and after balancing is implemented improve as the initial excess increases, although the relative difference is negligible. Before balancing is implemented, the fill rate is slightly higher than under balancing until the excess quantity is greater than the duration of the simulation. This gap widens as the quantity increases until all demand can be met at Patuxent River (*high demand*). When balancing is enabled, it actually improves the overall fill rate for the few instances in which Oceana's MSL is not zero, preventing the occurrence of stock-outs.

As we increase the initial amount of excess, Subfigure 3.4b shows that the overall inventory management costs decrease because the more materiel a location holds, the less it needs to be procured. In this study, we are more interested in differences in costs before and after the implementation of balancing, not actual cost values. The difference in cost is relatively small as the amount of excess increases, even though the actual cost trend is moving downward. However, there are two exceptions to this: 180 days of excess and 720 days of excess. With 180 days of excess, balancing is more expensive, but with 720 days of excess, the reverse is true. We see an increase in cost with 180 days of excess under balancing because Patuxent River's on-hand inventory reaches its reorder point more times because it transfers some of its materiel to other locations, thus requiring less procurement. The mean and standard deviation of the MSL are 124.5 and 585.6, respectively. For example, if the MSL is 124, then the reorder point is 62. Without balancing, the on-hand inventory reaches the ROP 10 times, and with balancing, it reaches the ROP 11 times. Therefore, under balancing, an additional 62 items are procured, increasing inventory management costs. In general, differences in costs before and after balancing is implemented are minimal for all amounts of excess residing at Patuxent River (*high demand*).

In our original analysis, Oceana (*moderate demand*) has an MSL that is equal to zero, which causes fill rates to decrease under balancing compared to not balancing. This is due to the




Figure 3.4: Bar chart of fill rates and inventory costs before and after balancing for various amounts of initial excess at Patuxent River.

excess fulfillment of demand after a shortage has already occurred. We found that increasing the allowance at Oceana from zero to two reverses this trend. Now, we will see whether this insight holds for different amounts of initial excess. Although difficult to see in Subfigure 3.5a, the fill rate increases after the implementation of balancing for all quantities of excess except 90 and 180 days.

Subfigure 3.5b shows inventory management costs before and after balancing implementation. Except in the case of 180 days of excess, costs are lower after balancing. Furthermore, Subfigure 3.5b shows that inventory management costs decrease with the addition of an allowance at Oceana both before and after the implementation of balancing, compared to Subfigure 3.4b, which does not have an allowance.

In the two previous tables, excess always resides at Patuxent River (*high demand*). In general, the benefits of balancing are minimal or slightly negative when excess resides in the high-demand location. For the WIRE, ROPE item, we rerun the simulation, varying the amount of the initial excess at Oceana. As in Figure 3.5, we increase the allowance from zero to two and vary the initial excess from zero to 720 days. The results are shown in Figure 3.6. In Subfigure 3.6a, the fill rates before balancing stop increasing after 90 days of excess. This is because 90 days of excess fulfills all demand at Oceana. The implementation of balancing decreases fill rates for 90,




Figure 3.5: Bar chart of fill rates and inventory costs before and after balancing for various amounts of initial excess with an increase in allowance at Oceana. Excess resides at Patuxent River.

180, and 360 days of excess, but with 720 days of excess, Oceana can still meet its own demand while shipping excess to Patuxent River (*high demand*). With 720 days of excess, the fill rate under balancing is larger than that without balancing.

Under balancing, inventory costs decrease as the amount of excess increases. At 90 days, the inventory costs without balancing increase, while the costs with balancing decrease. This is because once the excess fulfills all the demands at Oceana, the addition of any more excess without balancing only increases holding costs. The difference between the inventory costs before and after balancing increases, making balancing more desirable as the amount of excess increases.

3.2.2 Alternative Location: Kings Bay

The transshipment literature finds that the benefits of balancing increase with commonality of demand among locations. Since the WSS proposed business rules allow for balancing between locations that support different platforms, we replace Oceana with Kings Bay, a ballistic missile submarine base. We do this to determine if there is reason to include these types of locations within the same balancing pool. Again, we must assess the commonality of demand between the three locations. Without common demand, transshipments among the three locations would have no impact, either positive or negative. If each location's stock is unique relative to that of the




Figure 3.6: Bar chart of fill rates and inventory costs before and after balancing for various amounts of initial excess at Oceana. This includes an allowance increase at Oceana.

others, then no transshipments are possible. The following analysis examines the commonality in demand for line items at each location. Figure 3.7 shows this relationship between the three locations using a Venn diagram.

Each circle represents a different location. The sum of the numbers within the interior of an individual circle represents that location's demand for unique items, while the numbers outside that circle represent inventory items that are not demanded at that location. The overlapping areas, or intersections, represent common demand for unique items at each of the locations. The intersection of all three circles represents commonly demanded items among all three locations. We select only the top five items, according to quarterly dollar value, for examination. Table 3.7 shows the number of common line items that fall into each of the unit price thresholds. The unit prices of two common items are over \$5,000, while the majority of the common items have unit prices of less than \$25.00. No item over \$5,000 is common to the submarine base and the NASs. Two of the resulting five items cost less than \$25.00 and are therefore ineligible for transshipment under the proposed balancing policy.

Figure 3.7 shows that 296 line items are commonly demanded by at least two of the three locations. This study will only examine the top five line items out of these 296, according to quarterly dollar value (demand $\times$ unit price). Table 3.7 shows the number of common line items that fall




Figure 3.7: Venn diagram of demand commonality at Lemoore, Kings Bay, and Patuxent River. (After Liskow 2013)

into each of the unit price thresholds.

Table 3.7: Common line items that fall within each of the proposed unit price thresholds. (After Liskow 2013)

	Kings Bay	Lemoore	Patuxent River
\$ 5,000 and above	0	2	2
\$ 75.00 to \$ 4,999.99	8	13	18
\$ 25.00 to \$ 74.99	14	11	22
< \$ 25.00	204	62	246
Total	226	88	288

* Note 10 items are carried at all three locations

Like the results seen with the three NASs, the new set of locations including Kings Bay generates lower relevant costs when the excess resides at the location with zero demand for a given item. Again, this is because, without balancing, the excess incurs a constant holding cost at the location with zero demand and the locations with demand must procure the same materiel. However, when the excess resides in the high-demand location, the impact is small and, in some cases, negative. Table 3.9 shows the percent change from implementing balancing for each of the five items when the excess resides at the high-, moderate-, and low-demand locations.

As in our original analysis for the three NASs, we address the observed departures from the

Table 3.8: Summary data for the five most commonly demanded items by quarterly dollar volume.

NIIN	Weight (lbs)	Cube (ft ³)	Nomenclature	Quarterly Dollar Volume		
				Unit Price	Qty	Total
014617380	10.00	0.1481	HOOK POINT,ARRESTING	\$ 8676.59	7	\$ 60,736.13
011696129	0.57	0.0729	HOSE,OXYGEN GENERATOR	\$ 944.66	6	\$ 5,667.96
011510838	2.36	0.1334	VALVE,LINEAR,DIRECTIONAL CONTROL	\$ 5,737.87	1.33	\$ 4,303.40
012030382	0.32	0.0027	CABLE,SPECIAL PURPOSE,ELECTRICAL	\$ 1.65	2211.75	\$ 3,649.39
011223338	0.50	0.0034	CLAMP,HOSE	\$ 10.55	315.1	\$ 3,328.53

* Quarterly dollar value is the highest value of the three sites

anticipated benefits of transshipment, specifically focusing on how Kings Bay's relatively high allowance impacts our results for two of the items under consideration. Since Kings Bay has a much higher allowance for the CABLE, SPECIAL and CLAMP, HOSE items, even though the NASs have a higher demand for these items, the fill rate decreases under balancing. The holding costs also increase under balancing, which did not occur in the original analysis. In the following section, we will explore how the relatively high allowance compared to demand negatively impacts the anticipated benefits of balancing.

3.2.3 Alternative Location: Fill Rate Decreases

In Section 3.1.1, we observed that an anticipated benefit of balancing is an improvement in the fill rate. However, the fill rate decreases for the HOSE, OXYGEN and the VALVE, LINEAR under balancing for the same reason as for the WIRE, ROPE and CONE AND ROLLER items found in the original analysis. Due to one location's low demand, it has an MSL equal to zero for the majority of the simulation. Excess materiel fulfills demand but does not prevent shortages. The only difference in this analysis for the HOSE, OXYGEN and the VALVE, LINEAR items occurs when the initial excess resides at the zero-demand location. Unlike the case of the WIRE, ROPE and the CONE AND ROLLER items, the fill rate actually improves under balancing when excess resides in the low-demand location. This is because the improvement in the fill rate at the high-demand location impacts the overall fill rate more than at the moderate-demand location. As in our original analysis, increasing the allowance from zero to two once again improves the fill rate. By examining the alternative location of Kings Bay, we discover that the fill rate for the CABLE, SPECIAL and CLAMP, HOSE items decreases for a reason that is specific to this location and, therefore, not discussed in the original analysis.

Table 3.9: Cost and service impact of inventory balancing. Percent change represents the difference between the fill rate and cost values with and without balancing.

Positive percentage indicates improvement, while negative percentage indicates decrease in the value after the implementation of balancing.

Location of Excess <i>Item Description</i>	% Improvement		Daily
	Fill Rate	Total Cost (%)	Total Cost (\$)
Location of Initial Excess			
High Demand Location			
HOOK POINT	0.00 %	2.54 %	11.96
HOSE, OXYGEN	-0.05 %	1.39 %	0.76
VALVE, LINEAR	-6.50 %	13.61 %	5.29
CABLE, SPECIAL	-0.01 %	-2.39 %	-1.12
CLAMP, HOSE	-0.02 %	-4.61 %	-0.65
Location of Initial Excess			
Medium Demand Location			
HOOK POINT	0.00 %	2.52 %	12.23
HOSE, OXYGEN	-1.53 %	8.50 %	5.00
VALVE, LINEAR	-6.95 %	22.38 %	9.68
CABLE, SPECIAL	-0.01 %	3.86 %	2.84
CLAMP, HOSE	-0.02 %	14.80 %	3.03
Location of Initial Excess			
Low or No Demand Location			
HOOK POINT	0.00 %	13.30 %	70.83
HOSE, OXYGEN	0.30 %	11.92 %	7.29
VALVE, LINEAR	0.16 %	32.35 %	16.05
CABLE, SPECIAL	-0.01 %	4.42 %	2.46
CLAMP, HOSE	-0.02 %	14.62 %	2.99

3.2.4 How a High Allowance Relative to Demand Negatively Impacts Fill Rates under Balancing

In this section, we examine how fill rates decrease under balancing due to the comparably high allowance at Kings Bay (*moderate demand*) for items with less demand than at the high-demand location. Table 3.10 shows the allowance at the three locations and the associated two-year demand. Considering Kings Bay (*moderate demand*), notice that the allowance for the CABLE, SPECIAL item is slightly higher than its two-year demand and that the inventory level for the CLAMP, HOSE item will never breach the ROP in a two-year period. However, Patuxent River (*high demand*) has a greater demand for these items but an allowance of zero. The differences in the allowances at the NASs and submarine community could be attributable to the fact that these parts are less critical to the aviation community.

Table 3.10: Allowances and two-year demand for the CABLE, SPECIAL and CLAMP, HOSE items.

<i>Item</i>	<i>Allowance</i>	<i>Two Year Demand</i>
CABLE, SPECIAL, ELECTRICAL		
Patuxent River	0	7,647
Lemoore	0	0
Kings Bay	180	192
CLAMP, HOSE		
Patuxent River	0	640
Lemoore	0	0
Kings Bay	200	70

Figure 3.8 shows the relationship between Patuxent River (*high demand*) and Kings Bay (*moderate demand*) for these two items. As previously discussed, preventive transshipments are only beneficial when materiel moves to a location that is more likely to experience a shortage. However, Kings Bay (*moderate demand*) is less likely to experience shortages due to its high allowance and low demand. As in the case of the VALVE, REGULATING item in the original analysis, Patuxent River's demand for CABLE, SPECIAL and CLAMP, HOSE items is erratic, which sets the MSL low when future demand will be high and vice versa. Again, this is due to the level setting frequency and the demand pattern for these items. The preventive transshipment policy within Navy ERP uses the relationship between the on-hand inventory and the MSL to determine where the next shortage is likely to occur. However, at Kings Bay, this relationship is a poor indicator of where the next stock-out will occur. For a majority of the simulation, the on-hand inventory at Kings Bay (*moderate demand*) will be below its MSL but above its ROP. Excess materiel at Patuxent River (*high demand*) is shipped to Kings Bay, reducing its ability to prevent future stock-outs. The CABLE, SPECIAL, ELECTRICAL and CLAMP, HOSE items illustrate this behavior.

CABLE, SPECIAL, ELECTRICAL (NIIN 012030382)

This item is the least expensive, with a unit price of \$1.65. Under the WSS proposed balancing rules, it is ineligible for transshipment. The demand for the CABLE, SPECIAL, ELECTRICAL item at Patuxent River over the two-year period has a mean of 7,647.05 and a standard deviation of 121.8. At Kings Bay, it has a mean of 192.1 and a standard deviation of 13.7. Kings Bay also has an allowance of 180, which is not exceeded in 20.2 percent of the 2,000 replications.

When excess materiel resides at the high-demand location of Patuxent River, balancing decreases the fill rate by 1.08 percent. The fill rate at Kings Bay (*moderate demand*) is 100 percent without




Figure 3.8: Inventory balancing within Navy ERP screens locations for materiel above the MSL to be redistributed to locations with inventory below the MSL. When one location's allowance is much higher than its demand, it will continue to receive excess from other locations, even though it is probably the location least likely to experience a shortage.

balancing due to its high allowance quantity. When Patuxent River ships its excess materiel to Kings Bay (*moderate demand*), the materiel is not available to prevent future shortages at its origin. The fill rate at Kings Bay (*moderate demand*) is 100 percent, regardless of whether balancing is employed. Even though Patuxent River (*high demand*) does not receive excess at the beginning of the simulation, its MSL could decrease between the 90-day level setting, thus producing excess. Patuxent River ships this resulting excess to Kings Bay (*moderate demand*) even when the initial excess materiel resided in the low-demand location of Lemoore.

The no threshold policy allows for transshipment among all locations regardless of unit price. Since the allowance at Kings Bay is much higher than its demand, it will continue to receive excess from Patuxent River (*high demand*), even though it is probably the location that is the least likely to experience a shortage. When the inventory at Kings Bay (*moderate demand*) falls below the MSL, it is not the best indication of an impending stock-out. The CABLE, SPECIAL, ELECTRICAL item in Table 3.9 illustrates this situation. When Patuxent River (*high demand*) ships its excess materiel to Kings Bay (*moderate demand*), this materiel is not available to prevent future shortages at its origin or improve the fill rate at its destination. Instead, a better indication of an impending stock-out is when the on-hand inventory falls below its ROP. This alternative policy, based on the reorder point, will prevent Patuxent River (*high demand*) from transshipping materiel to Kings Bay (*moderate demand*), with its already high allowance. Table 3.11 shows the fill rates for the policy without balancing, the no-threshold policy, and the ROP alternative

balancing policy.

Table 3.11: Fill rate for NIIN 012030382 for each location after running the simulation 2,000 times with 30 days of excess, including an alternative balancing policy.

<i>Balancing</i>	<i>Location</i>	<i>Initial Excess</i>	<i>Fill Rate (%)</i>	
			μ	σ
<i>NLS</i>	Patuxent River	<i>Yes</i>	41.90%	28.62%
	Lemoore	<i>No</i>	100.00%	0.00%
	Kings Bay	<i>No</i>	100.00%	0.00%
<i>No threshold</i>	Patuxent River	<i>Yes</i>	40.82%	29.13%
	Lemoore	<i>No</i>	100.00%	0.00%
	Kings Bay	<i>No</i>	100.00%	0.00%
<i>Reorder point</i>	Patuxent River	<i>Yes</i>	41.12%	29.00%
	Lemoore	<i>No</i>	100.00%	0.00%
	Kings Bay	<i>No</i>	100.00%	0.00%

The alternative reorder balancing policy improves system-wide fill rates by 0.12 percent compared to the no-threshold balancing policy, but the system-wide fill rate decreases by 0.32 percent compared to the no-balancing policy. The alternative policy also slightly improves procurement costs compared to the no-threshold balancing policy. However, the total cost of the alternative policy is higher than that of the no-balancing policy, which is mostly due to a 2.69 percent increase in backorder costs.

Again, the differences in the fill rates for Patuxent River (*high demand*) among the three policies seem minor, as shown in Table 3.11. However, the performance of a matched pair *t* test shows that there is sufficient evidence to conclude that the differences in fill rates after balancing are statistically significant.

CLAMP, HOSE (NIIN 011223338)

This item is the second least-expensive, with a unit cost of \$10.56. Under the WSS proposed balancing rules, it is not eligible for transshipment. The demand for the CLAMP, HOSE item over the two-year period at Patuxent River (*high demand*) has a mean of 640.3 and a standard deviation of 121.8. At Kings Bay (*moderate demand*), it has a of 20.0 and a standard deviation of 5.0. Kings Bay (*moderate demand*) also has an allowance of 200 and its reorder point is never reached in the two-year period. Table 3.12 shows the fill rates for this case.

Table 3.12: Fill rates for NIIN 011223338 for each location after running the simulation 2,000 times with 30 days of excess.

<i>Balancing</i>	<i>Location</i>	<i>Initial Excess</i>	<i>Fill Rate (%)</i>	
			μ	σ
<i>NLS</i>	Patuxent River	<i>Yes</i>	53.85%	29.94%
	Lemoore	<i>No</i>	100.00%	0.00%
	Kings Bay	<i>No</i>	100.00%	0.00%
<i>No threshold</i>	Patuxent River	<i>Yes</i>	51.96%	21.47%
	Lemoore	<i>No</i>	100.00%	0.00%
	Kings Bay	<i>No</i>	100.00%	0.00%

As in the case of the previous item, CABLE, SPECIAL, ELECTRICAL, only the fill rate at the high-demand location of Patuxent River decreases. This is because the allowance for this item at Kings Bay (*moderate demand*) is 200, and the maximum two-year demand for 2,000 replications is 43. Without balancing, Patuxent River (*high demand*) prevents the occurrence of stock-outs through its excess materiel. However, under balancing, its excess is transferred to Kings Bay (*moderate demand*), whose allowance ensures that it will maintain a 100 percent fill rate with or without excess from Patuxent River (*high demand*). For two of the five items, Kings Bay (*moderate demand*) maintains an allowance above the two-year historical demand total, whereas Patuxent River does not have an allowance for the item.

3.2.5 Alternative Location: Holding Costs Increase

Unlike in our original analysis, in which holding costs never increased under balancing, holding costs increase for the CLAMP, HOSE item when the excess resides at the high-demand location. This is due to the comparatively high allowance for these items at Kings Bay (*moderate demand*) for these items. The maximum demand at Kings Bay (*moderate demand*) will not decrease the on-hand inventory below its ROP. Before balancing, the high allowance allows Kings Bay (*moderate demand*) to meet demand in all 2,000 replications, without ever having to replenish its stock. In contrast, Patuxent River sets its MSL based on demand every 90 days. Its fill rate is a little over 50 percent due to its erratic demand. When the Patuxent River (*high demand*) ships materiel to Kings Bay (*moderate demand*), it decreases its pool of excess, but also increases the on-hand stock at Kings Bay. The redistribution of materiel within the system, thus increases holding costs for two reasons. First, any materiel shipped to Kings Bay (*moderate demand*), with its high allowance and low demand, only increases the amount of materiel on its shelves.

Second, the surplus of materiel does not provide any additional benefits, such as an increase in the fill rate and decrease in stock-out costs, since Kings Bay’s allowance is more than enough to fulfill all demand over the two-year period.

3.2.6 Resetting the Allowance

In the previous sections, we discovered that the reason that the fill rate decreased and holding costs increased was the comparably high allowance at Kings Bay (*moderate demand*) for items with less demand than at Patuxent River (*high demand*). In this section, we adjust the allowance at Kings Bay (*moderate demand*) and Patuxent River (*high demand*) based on the observed mean MSL for the CABLE, SPECIAL, ELECTRICAL and CLAMP, HOSE items from the simulation. Remember, previous demand does not determine an item’s allowance. Instead, allowances are based on the item’s importance to the weapon system. Therefore, the disparity between Patuxent River (*high demand*) and Kings Bay (*moderate demand*) may depend on the weapon platforms supported at these locations.

The alternative allowances in Table 3.13 are based on the simulation’s mean and standard deviation for the MSLs when no allowance is set. Since allowances are based on more than demand, we assume that if Kings Bay (*moderate demand*) qualifies for an allowance, then so should Patuxent River (*high demand*). Due to the high demand for the CABLE, SPECIAL, ELECTRICAL item at Patuxent River, its overall allowance is greater than the original allowance assigned to Kings Bay (*moderate demand*); however, the allowance at Kings Bay (*moderate demand*) is reduced by 92 percent. The overall allowance for the second item, CLAMP, HOSE, is reduced by 78 percent, and the allowance at Kings Bay (*moderate demand*) is reduced by 98.5 percent.

Table 3.13: Original and alternative allowances.

<i>Item</i>	<i>Original Allowance</i>	<i>Alternative Allowance</i>
CABLE, SPECIAL, ELECTRICAL		
Patuxent River	0	250
Lemoore	0	0
Kings Bay	180	14
CLAMP, HOSE		
Patuxent River	0	40
Lemoore	0	0
Kings Bay	200	3

Table 3.14 demonstrates the contrast between the original and alternative allowances before and

after the implementation of balancing. For the first item, CABLE, SPECIAL, ELECTRICAL, balancing still increases total costs when the excess resides in the high-demand location, but by much less than when using the original allowance. Even though the fill rate remains the same before and after balancing, backorder costs increase. This is due to how stock-out occurrences are combined. Regardless of how little a shipment weighs, there is a minimum cost to ship material. For this item, the costs of shipping one or 11 items are identical. Although Patuxent River (*high demand*) experiences the same amount of stock-out occurrences under balancing, they tend to become more spread out. This prevents expedited stock-out orders from being combined, resulting in higher backorder costs. Therefore, although the fill rate stays the same under balancing, backorder costs increase for this item.

Table 3.14: Cost and service impact from inventory balancing. **Positive percentage indicates improvement, while negative percentage indicates decline in the value after implementing balancing.**

Item Description <i>Location of Excess</i>	% Improvement	
	<i>Fill Rate</i>	<i>Total Cost</i>
CABLE, SPECIAL, ELECTRICAL		
Original Allowance		
High Demand Location	-0.01%	-2.39 %
Medium Demand Location	-0.01%	3.86 %
Low Demand Location	-0.01 %	4.42 %
Alternative Allowance		
High Demand Location	0.00 %	-0.95 %
Medium Demand Location	2.61%	16.61 %
Low Demand Location	17.52%	17.86%
CLAMP, HOSE		
Original Allowance		
High Demand Location	-0.02%	-4.61 %
Medium Demand Location	-0.02%	14.80%
Low Demand Location	-0.02%	14.62 %
Alternative Allowance		
High Demand Location	0.16 %	0.03 %
Medium Demand Location	0.31%	24.16 %
Low Demand Location	0.31%	26.92%

Table 3.14 shows the percent change before and after the implementation of balancing. Table 3.15 shows the fill rates for the original and alternative allowances for the CLAMP, HOSE item when demand resides at Kings Bay (*moderate demand*). Although the fill rate at Kings Bay (*moderate demand*) decreases when the allowance is decreased from 180 to 3, this decrease is small. The overall allowance quantity among the three locations actually decreases from 200 to 43 overall. Due to its erratic demand, the fill rate at Patuxent River (*high demand*) increases considerably when an allowance is incorporated.

Table 3.15: Daily on-hand inventory and procurements for each location for NIIN 011223338 after running the simulation 2,000 times with 30 days.

<i>Balancing</i>	<i>Location</i>	<i>Initial Excess</i>	<i>Original Fill Rate (%)</i>		<i>Alternative Fill Rate (%)</i>	
			μ	σ	μ	σ
<i>NLS</i>	Patuxent River	<i>No</i>	38.36%	26.92%	95.73%	3.61%
	Lemoore	<i>No</i>	100.00%	0.00%	100.00%	0.00%
	Kings Bay	<i>Yes</i>	100.00%	0.00%	100.00%	0.00%
<i>No threshold</i>	Patuxent River	<i>No</i>	36.53%	21.47%	96.20%	3.64%
	Lemoore	<i>No</i>	100.00%	0.00%	100.00%	0.00%
	Kings Bay	<i>Yes</i>	100.00%	0.00%	96.07%	5.39%

3.2.7 Transshipment Costs Still Have a Minimal Impact

As in our original analysis, transshipment contributes little to overall inventory management costs. Table 3.16 shows that transshipment represents less than three percent of total costs in all 15 scenarios. This reinforces the original observation that transshipment costs contribute little to overall inventory management costs.

Table 3.16: Percent of total cost attributed to transshipments.

<i>ITEM</i>	<i>% of Total Cost</i>	<i>ITEM</i>	<i>% of Total Cost</i>
HOOK POINT, ARRESTING		HOSE, OXYGEN GENERATOR	
High Demand Location	0.004%	High Demand Location	0.007%
Medium Demand Location	0.011%	Medium Demand Location	0.0042%
Low Demand Location	0.015%	Low Demand Location	0.043%
VALVE, LINEAR, DIRECTIONAL		CABLE, SPECIAL, ELECTRICAL	
High Demand Location	0.015%	High Demand Location	0.537%
Medium Demand Location	0.023%	Medium Demand Location	2.450%
Low Demand Location	0.029%	Low Demand Location	1.714%
CLAMP, HOSE			
High Demand Location	0.548%		
Medium Demand Location	1.342%		
Low Demand Location	1.547%		

3.3 Summary

In order to evaluate the three different transshipment policies, we run a simulation to determine whether balancing improves fill rates and costs compared to not balancing. Our results indicate that the savings from reductions in procurement drive the transshipment decision. When excess is available to more locations, the system procures less materiel, increasing the likelihood that inventory balancing will lead to lower overall costs. We find that backorder costs are the next most influential factor in determining whether balancing generates savings. In our analysis, the difference in backorder costs before and after the implementation of balancing is minimal when

the excess resides in a location with a relatively high demand compared to the amount of excess. However, the difference becomes considerable as the amount of excess becomes relatively larger than demand, growing from a difference of only a few cents to dollars. This explains why inventory balancing is more beneficial when the excess resides in the moderate- and low-demand locations. When we consider the totality of inventory management costs, holding costs make up the largest proportion in some cases. However, there is relatively little difference in holding costs between balancing and not balancing compared to differences in procurement and backorder costs, except when the excess resides in a location with no demand. In this case, the difference between balancing and not balancing is in the hundreds of dollars. Finally, we discover that transshipment costs only exist under balancing. However, as a percentage of total costs, this cost represents less than three percent in every tested case.

When balancing performs poorly, all the excess resides in the high-demand location. The two-year savings for the nine unique examined items is \$5,322 under balancing. When the excess resides in the low-demand location, the savings for the nine items is \$15,559,572 under balancing. However, the placement of all of the excess in either the high- or low-demand locations, as in our simulation, probably does not represent the true distribution of materiel. The excess more likely resides in some combination of these two extremes. We do know that millions of dollars in excess is available at Lemoore, Oceana, and Patuxent River and could be used to offset procurement if balancing were implemented, as described in Section 2.1.1.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 4:

Conclusion and Recommendations

The Navy Enterprise Resource Planning (ERP) Single Supply Solution can provide many potential benefits through the centralization of inventory management into a single database. The focus of this study centers around one of these benefits: inventory balancing. Previously impossible with Navy legacy inventory management systems, the implementation of inventory balancing requires a connected information system, which is now possible through Navy ERP. This solution now provides unprecedented visibility for Navy-owned materiel at ashore activities. This visibility subsequently enables excess materiel to be redistributed in order to prevent the unnecessary procurement of Navy and Defense Logistic Agency (DLA) materiel. Finally, it is possible to reduce inventory management costs and improve service through inventory balancing.

In the first study to analyze inventory balancing within the Navy, we measure the impact of the business rules that govern balancing. Our approach compares fill rates and inventory management costs before and after the implementation of balancing policies, enabling decision makers to determine which are the most cost-effective. This thesis is intended to serve as a baseline for future studies. The goal of our analysis is to measure the benefits of the Weapon Systems Support (WSS) proposed balancing policy and propose modifications that would enhance its strengths and minimize its weaknesses. We examine DLA and General Service Administration (GSA) materiel funded by Budget Project 28 (BP-28). Currently, this materiel is held at 39 different locations around the world. Although this analysis examines BP-28 materiel, the resulting insights can be applied to inventory balancing in general and are applicable to other types of materiel.

4.1 Discussion

Inventory balancing does not increase fill rates and does not decrease inventory management costs in all situations. In fact, the fill rate actually decreases under balancing more often than it improves in the examined instances. In the majority of cases, it decreases by less than one percent. However, balancing does decrease excess materiel. The total two-year mean inventory costs for the nine unique examined items generates savings of an average of \$591 per item compared to not balancing when the excess resides at the high-demand location. Note that this average includes items that actually result in higher costs. We only include nine unique items in the study because one item is common in both our original naval air station (NAS)-only

analysis and the alternative analysis, which includes a submarine base. The mean savings per item of \$591 is based on the worst observed case. When the excess resides in the moderate- and low-to-no-demand locations, the savings are even greater. We find that expansion in the number of items and locations provides a much greater potential for savings.

The following outlines the main insights of our analysis. In Section 4.2, we provide recommendations for the modification of business rules based on these insights.

Dollar Volume versus Unit Price. The WSS proposed balancing policy uses unit price to determine geographic eligibility for balancing. Our analysis finds that 10 percent of the items at the three NASs contribute between 78 to 96 percent of inventory costs. Some of these influential items cost less than \$25, which is the required threshold for inventory balancing eligibility under the WSS proposed balancing rules. Balancing reduces inventory management costs for three items in this study that cost less than \$25.

When Excess Is Relatively Small Compared to Demand. For four of the five original items and for all items in our alternative analysis, which includes a submarine base, balancing represents only a small improvement, and in some instances, costs increase when excess resides in the high-demand location compared to when the excess resides at either of the other locations. The only exception to these observations occurs with the HOOK POINT item in our original analysis. This is the only instance when there is demand for the materiel at all three locations and demands for the high- and moderate-demand locations are very similar. Without balancing, excess at the high-demand location will not remain in inventory very long. Balancing will incur a transportation cost and will potentially increase stock-out occurrences at the high-demand location because the inventory has been moved to the lower-demand locations.

Erratic Demand. Some items experience consecutive months with zero demand, followed by a month or two with very high demand. This results in a low maximum stock level (MSL) when future demand is high. In this situation, inventory balancing transfers materiel away from a location that will soon experience high demand to a location that is less likely to experience demand. When we pair a location with this erratic demand pattern to a location with a high allowance that exceeds the two-year mean demand, the negative impact of balancing is only exacerbated.

MSL Equal to Zero. Inventory balancing screens locations for materiel above the MSL to

be redistributed to locations with inventory below the MSL. Because inventory balancing moves excess to a location with an MSL greater than zero, the occurrence of stock-outs will most likely be prevented. However, if the MSL is set to zero, this transfer will not prevent stock-outs but will prevent procurement and reduce holding costs.

Churn and Transshipment Costs. In one or two of the 2,000 replications for the ten examined items, materiel is transferred back and forth between locations once or twice. Fill rates, inventory management costs, and materiel availability are not impacted by churn. Transshipment costs make up the smallest portion of inventory management costs. In 87 percent of all the scenarios, transshipment contributes less than one percent to inventory management costs. In no cases did it contribute to more than three percent.

More Excess. In our analysis, we see little benefit from balancing when excess is small and demand is relatively high. However, when the amount of excess greatly exceeds the demand at a location, the benefits of balancing only increase as excess increases.

High Allowance Relative to Demand. Navy ERP uses the relationship between the on-hand inventory and the MSL to determine where the next shortage is likely to occur. However, this relationship is a poor indicator of where the next stock-out will occur when the allowance is more than twice the two-year mean demand. The on-hand inventory in this instance is below its MSL but above its ROP. Other locations ship their excess materiel to this type of location, increasing the likelihood that future stock-outs will occur. The location receiving the inventory obtains no benefits and its holding costs increases because it already has enough inventory to meet its demand. Inventory balancing decreases the fill rates and increases the holding costs of the system when the allowance is relatively high compared to demand.

4.1.1 Research Questions Answers

This study examines six research questions to determine the benefits of competing transshipment policies:

1. Do fill rates at individual retail sites improve when a preventive lateral transshipment policy (inventory balancing) is implemented?
 - In some case, fill rates modestly improve with balancing.
 - Fill rates may decrease at both individual retail sites and over the entire system.

- If fill rates decrease, they often decrease by less than one percent.
2. Given fill rate decreases, how do the different demand distributions affect the benefits of the preventive transshipment policy?
 - A low, infrequent demand distribution leads to an MSL equal to zero. If the MSL at a location is set to zero, excess transfers to it will not prevent stock-outs from occurring, but they will prevent procurement and reduce holding costs.
 - Lumpy demand also impacts balancing by setting the MSL low when future demand may be high. In situations where the MSL is low when demand is high, inventory balancing transfers materiel away from a location that will soon experience high demand to a location that is less likely to experience demand.
 3. Regardless of whether there are improvements, how does redistribution affect procurement?
 - The generation of savings from less procurement is the main reason inventory management costs decrease under balancing.
 - With excess available through balancing, locations procure less materiel.
 4. Do different balancing criteria at the retail level result in differing fill rates between the retail establishments? If the proposed balancing criteria fail to improve on those policies without balancing, do alternative criteria (i.e., a different set of balancing business rules) provide benefits?
 - Although the system-wide fill rate decreases, balancing usually improves the fill rate of at least one location. The only exception is when one of the location's allowances is so high that it achieves a 100 percent fill rate without balancing.
 - Balancing leads to improvements when the proposed business rules are modified to use the reorder point (ROP) instead of the MSL to determine which locations will receive excess.
 - Conducting level setting more frequently also improves balancing performance for the proposed business rules.
 5. If any of the sets of balancing decision rules tend to improve fill rates, does one set of balancing criteria tend to outperform the others? If so, which alternative reduces procurement costs by more than it increases transportation costs?
 - Inventory balancing does not increase fill rates or decrease inventory management costs in all situations.
 - Using the ROP instead of the MSL mitigates the possible negative impacts of balancing on fill rates, but also reduces the savings generated from the offsetting of procurement through excess.

6. Assess which policies tend to maintain lower inventories over time. What is the impact of preventive transshipment on the amount of excess and stock above the MSL as compared to the case with no inventory balancing?
 - Inventory balancing effectively minimizes excess materiel using the proposed business rules.
 - However, modifying the proposed business rules (e.g., using dollar volume, using ROP instead of MSL, etc.) can mitigate the potentially negative impact of balancing on fill rates and inventory management costs while still allowing for many of the benefits to be realized.

4.1.2 Lessons Learned

In our study, we find that of the three balancing policies under consideration, *no lateral shipments (NLS)*, *proposed*, and *no threshold*, not a single one clearly improved both fill rates and inventory management costs in all situations. However, our analysis uncovered key insights to drive transshipment decisions regarding whether to balance or not. The following circumstances delineate when balancing is preferable to not balancing.

1. Balancing improves fill rates and costs when the amount of excess is greater than the future demand.
 - The impact of balancing is the greatest when the excess resides in the location with zero demand. This reduces holding and procurement costs.
 - When the amount of excess is less than the mean 90-day demand, the impact from balancing is minimal and, in some cases, detrimental.
2. When a location's on-hand inventory is zero, the shipment of excess to this location does not improve fill rates but can reduce inventory management costs.
 - Raising the MSL at a location by increasing its allowance from zero to two can reverse this trend.
 - Increasing the allowance from zero to two can also improve inventory management costs by reducing backorder costs.
3. Erratic demand results in setting the MSL low when future demand is high. Inventory balancing exacerbates this issue by shipping excess away from where demand is more likely to occur, thus decreasing the fill rate, which increases backorder costs.
 - Increasing the frequency of level setting reverses this trend.
 - Performing inventory balancing only after level setting but before procurement also minimizes this issue.

4. When a location has a relatively high allowance compared to demand (e.g., an allowance that is more than twice the mean two-year demand), the use of the MSL as a criterion for eligibility to receive excess makes balancing less preferable.
 - The use of the ROP instead of the MSL as the criterion for eligibility to receive excess reduces and, in some cases, eliminates the negative impact from balancing that results when one location has a high allowance relative to demand.
 - Resetting the allowance to a more reasonable level greatly reduces actual costs, regardless of whether balancing is employed. However, improvements from balancing are greater than the situation without balancing.
5. Transshipment costs contribute less than three percent to inventory management costs in all situations. In none of the cases did transshipment costs determine whether balancing was beneficial.

4.2 Recommendations to Improve the Balancing Policy

1. Use dollar volume instead of unit price to determine eligibility for balancing. Limit balancing to the 10 percent of items that account for the majority of inventory management costs. This will reduce the number of shipments between locations by only transshipping the items that will most likely result in reduced inventory costs.
2. Instead of balancing before the generation of purchase requests twice a week, only run inventory balancing after quarterly level setting, which occurs every 90 days. Although this will reduce the potential savings from inventory balancing, it will also reduce its potential risks. This recommendation will decrease the number of transfers between locations, ensure that materiel is packaged together, and decrease some of the observed negative impacts of inventory balancing. We find that increasing the frequency of level setting reverses the negative impact of balancing on costs and fill rates. Increasing the frequency of level setting may not be feasible, but limiting the occurrence of balancing to only after quarterly levels have been set results in similar benefits.
3. Remove balancing limitations between the geographic regions within the continental United States. The shipping times and prices between Oceana and Patuxent River are the same as between Oceana and Lemoore. From the inventory cost and materiel availability perspectives, these geographic limitations are not relevant.
4. Do not transfer materiel to locations with an on-hand inventory levels below their MSL. Instead, transfer inventory to locations with on-hand inventory levels below their ROP. This limits the potential negative impact of balancing, as seen in our analysis section, while

preserving most of its benefits.

5. Do not transfer excess materiel from locations if the amount of excess is less than the 90-day demand forecast. When the excess resides in a location in which demand is relatively high compared to the amount of excess, balancing generates minimal benefits and is, in some cases, detrimental.

4.3 Future Research

In this section, we provide suggestions for follow-on research topics borne out of the analysis outlined in this study.

More Locations. Budget Project 28 (BP 28) includes 39 locations. This study only considers three locations at a time. The inclusion of all locations within the continental United States would be a beneficial enhancement to the study. A common finding in the transshipment literature attests that inventory balancing will become more beneficial as the number of locations increases.

More Items. This study examines nine unique items. Although some general trends emerge from the analysis, looking at additional items will either support or provide further insights into the impact of balancing.

Tie Inventory Balancing Business Rules to Level Setting or Demand Forecasting. Navy ERP uses the relationship between the on-hand inventory and the MSL to determine where the next shortage is likely to occur. We have seen that the ability of inventory balancing to positively impact fill rates and inventory costs depends on level setting. Therefore, we suggest further research that ties balancing to level setting or demand forecasting. There is a symbiotic relationship between inventory balancing and level setting. For a majority of the analysis, level setting remains constant, but we did see that changes in the frequency of level setting greatly increase the benefits of balancing.

Allowancing. In our analysis, we find that allowancing can greatly influence fill rate and inventory management costs. The visibility provided by the ERP allows for allowancing to be observed across all the sites. Additional research on allowancing has the potential to improve fill rates and inventory management costs. In two cases within our study, we change the allowances and generate drastic improvement in fill rates and costs, regardless of whether balancing is employed.

Real World Data. Once inventory balancing is implemented, there will be the opportunity to analyze data on its real-world impact. This study uses a simulation to represent the real system. In order to determine whether the resulting benefits found in this study reflect real-world data, we suggest a data analysis project to measure the impact of inventory balancing on fill rates and costs. The simulation in our analysis shows that the implementation of balancing has the potential to generate negative outcomes. Data analysis may provide additional insights to improve balancing that were unattainable through the simulation.

Possible Extensions of the Simulation. Although the purpose of the simulation is to measure the performance of balancing business rules, it is well-suited to measuring the impact of other inventory management decisions. For example, it can also test how different allowance settings impact fill rates and inventory costs. The code is easily modified to test different level setting algorithms. The only difficulty here is in coding the alternative level setting procedures in Visual Basic for Applications. It also can be expanded to include more than three locations. In this case, the expansion includes the modification of historical demand inputs and the duplication of functions for the additional locations. The balancing function can balance up to all of the 39 BP 28 locations without requiring additional coding.

References

- Allen, S. C. 1958. Redistribution of total stock over several user locations. *Naval Research Logistics Quarterly* **5**(4) 337–345.
- Amaral, J., D. Kuettner. 2008. Analyzing supply chains at HP using spreadsheet models. *Interfaces* **38**(4) 228–240.
- Banks, J., S. Buckley, S. Jain, P. Lendermann, M. Manivannan. 2002. Panel session: Opportunities for simulation in supply chain management. E. Yücesan, C. H. Chen, J. L. Snowdon, J. M. Charnes, eds. *Proc. 2002 Winter Simulation Conf.* (WSC'02), Institute of Electrical and Electronics Engineers, Washington DC, 1652–1658.
- Cachon, G. P. 2001. Stock wars: Inventory competition in a two-echelon supply chain with multiple retailers. *Oper. Res.* **49**(5) 658–674.
- Cardillo, J. A. 2012, November, 30. Electronic message between author and John Cardillo, Operations Research, Naval Supply Systems Command Weapon Systems Support.
- Ekren, B. Y., S. S. Heragu. 2008. Simulation based optimization of multi-location transshipment problem with capacitated transportation. S. J. Mason, R. R. Hill, L. Mönch, O. Rose, T. Jefferson, J. W. Fowler, eds. *Proc. 2008 Winter Simulation Conf.* (WSC'08), Institute of Electrical and Electronics Engineers, Washington DC, 2632–2638.
- Jönsson, H., E. A. Silver. 1987. Analysis of a two-echelon inventory control system with complete redistribution. *Management Sci.* **33**(2) 215–227.
- Kukreja, A., C. P. Schmidt. 2005. A model for lumpy demand parts in a multi-location inventory system with transshipments. *Comput. Oper. Res.* **32**(8) 2059–2075.
- Lasater, J. T. 2013, July 3. Electronic message between author and LT Josh Lasater, Operational Supply Planning, Naval Supply Systems Command Weapon Systems Support.
- Law, A. M., W. D. Kelton. 2000. *Simulation Modeling and Analysis*, 3rd ed. McGraw-Hill, New York.
- Liskow, E. M. 2013, August 16. Electronic message between author and Eric Liskow, Operations Research, Naval Inventory Control Point.
- McGavin, E. J., J. E. Ward, L. B. Schwarz. 1997. Balancing retailer inventories. *Oper. Res.* **45**(6) 820–830.
- McNeeley, D. C. 2013, January 7. Electronic message between author and Donald McNeeley, Transportation and Distribution Optimization, Naval Supply Systems Command Global Logistics Support.

- Naval Supply Systems Command Weapon Systems Support. 2011. Single national inventory guide. Unpublished guide, Naval Supply Systems Command, Mechanicsburg, PA.
- Needham, P. M., P. T. Evers. 1998. The influence of individual cost factors on the use of emergency transshipments. *Transportation Res. E* **34**(2) 149–160.
- Office of the Assistant Secretary of the Navy, United States Department of the Navy. 2011. Navy working fund: Fiscal year (FY) 2012 budget estimates. Navy Supply Management.
- Palko, B. D. 2013, July 30. Electronic message between author and Brian Palko, Repairables Distribution, Naval Supply Systems Command Weapon Systems Support.
- Ross, S. M. 2006. *Introduction to Probability Models*, 10th ed. Academic Press, San Diego, CA.
- Rudi, N., S. Kapur, D. F. Pyke. 2001. A two-location inventory model with transshipment and local decision making. *Management Sci.* **47**(12) 1668–1680.
- Silver, E. A., H. Naseraldin, D. P. Bischak. 2008. Determining the reorder point and order-up-to-level in a periodic review system so as to achieve a desired fill rate and a desired average time between replenishments. *J. Oper. Res. Soc.* **60**(9) 1244–1253.
- Silver, E. A., D. F. Pyke, R. Peterson. 1998. *Inventory Management and Production Planning and Scheduling*, 3rd ed. John Wiley & Sons, New York.
- Spackman, M. 2004. Time discounting and of the cost of capital in government. *Fiscal Studies* **25**(4) 467–518.
- Tagaras, G. 1999. Pooling in multi-location periodic inventory distribution systems. *Omega* **27**(1) 39–59.
- Tyworth, J. E., L. O'Neill. 1997. Robustness of the normal approximation of lead-time demand in a distribution setting. *Naval Res. Logist.* **44**(2) 165–186.
- United States Government Accountability Office. 2005. DOD business systems modernization: Navy ERP adherence to best business practices critical to avoid past failures (GAO-05-858). United States Government Accountability Office, Washington DC.
- United States Government Accounting Office. 1995. Defense inventory: Opportunities to reduce warehouse space (NSIAD 95 64). United States Government Accounting Office, Washington DC.
- Wackerly, D. D., W. Mendenhall, R. L. Scheaffer. 2008. *Mathematical Statistics with Applications*, 7th ed. Thomson Brooks/Cole, Belmont, CA.
- Wong, H., G. J. van Houtum, D. Cattrysse, D. Van Oudheusden. 2006. Multi-item spare parts systems with lateral transshipments and waiting time constraints. *Eur. J. Oper. Res.* **171**(3) 1071–1093.

Initial Distribution List

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