

Logistics Management Institute

Optimizing Spares Support: The Aircraft Sustainability Model

AF501MR1



October 1996

DTIC QUALITY INSPECTED 4

F. Michael Slay
Tovey C. Bachman
Robert C. Kline
T.J. O'Malley
Frank L. Eichorn
Randall M. King

© Logistics Management Institute 1996
All Rights Reserved

Prepared pursuant to Department of Defense Contract DASW01-95-C-0019. The views expressed here are those of the Logistics Management Institute at the time of issue but not necessarily those of the Department of Defense. Permission to quote or reproduce any part except for government purposes must be obtained from the Logistics Management Institute.

Logistics Management Institute
2000 Corporate Ridge
McLean, Virginia 22102-7805

19970122 101

Optimizing Spares Support: The Aircraft Sustainability Model

Executive Summary

Modern inventory management looks as much to quick response to provide materiel when requested as it does to buying and holding in anticipation of demand. Particularly when inventoried items are low cost, stable in demand, and may be procured quickly, the practice is to carry as little inventory as possible. Variants of the "just-in-time" principle, including indefinite delivery, indefinite quantity contracts with vendors, and the use of commercial overnight delivery services, allow managers to treat the holding of inventory as a last resort.

Supporting advanced military equipment, however, poses unique problems that demand more sophisticated treatment. A modern warplane contains many components designed to be removed and replaced when they fail rather than to be repaired in place. Many of these components — brake assemblies, avionics units, and engine fuel controls — are expensive enough to warrant repair in their own right at the operating location, a military maintenance depot, or a commercial contractor. Spares of these reparable components are needed to keep the even-more-expensive aircraft ready to fly while the failed component is being repaired. The typical small fleet sizes and low operating tempo of military aircraft (or similar end items) lead to sporadic and unstable demand patterns; the specialized nature of many components leads to small market size and procurement lead-times that can stretch into years. These two factors limit the utility of some of the newer commercial practices and argue in favor of large inventories to ensure support of the end items. Yet, when unit costs for some components approach \$1 million, the necessity to keep inventories as small as possible is obvious.

Military materiel management, then, must use as much of modern commercial practice as practical so that spares can be repaired and shipped quickly. But these methods do not eliminate the need to forecast requirements and make the best inventory decisions possible. Certainly, the decision of what expensive, long lead-time component spares should be procured for the inventory will always be a critical one with expensive consequences.

The Aircraft Sustainability Model (ASM), developed by the Logistics Management Institute for the United States Air Force (USAF), is a mathematical statistical model that computes optimal spares mixes to support a wide range of

possible operating scenarios. In contrast to the typical practice, the ASM sizes the spares inventory based explicitly on desired weapon system readiness levels, such as aircraft availability — the percentage of the fleet ready to fly a mission — rather than supply-oriented measures, such as stock on the shelf or percentage of demands filled.

The ASM is used by the USAF to determine spares kits to support squadron deployments in wartime and has been specially enhanced for the initial provisioning process. It is in use by the USAF for initial provisioning for the F-22 Advanced Tactical Fighter and the E-8 Joint Surveillance, Target Attack Radar System (JSTARS); it has been proposed as a Department of Defense standard. The ASM uses the typical component data — demand rates, repair times, unit cost, and so on — in concert with any of a wide range of operating scenarios — number of aircraft, time phasing of aircraft procurement, and operating tempo. The ASM then uses a marginal analysis approach, ranking possible additions to the inventory in terms of their probable benefit to aircraft availability divided by their procurement cost. Spares that have the greatest benefit per dollar appear at the top of this “shopping list.” Accumulated costs and resulting aircraft availability are tracked as the shopping list is formed to provide a curve relating overall funding and projected availability.

The curve can be used by logistics planners to formulate budgets and allocate resources. On the component level, item managers can use the shopping list to determine detailed buy requirements consistent with those aggregate funding decisions. The mathematical algorithm ensures that the shopping list is optimal — given a funding level, buying down the list until that level is reached provides the highest aircraft availability rate; conversely, given an availability rate target, the shopping list provides the least cost mix to reach that target. While these projections are probabilistic in nature, and unavoidably subject to uncertainty, they make the best possible use of the data available.

The ASM resides on a PC platform with a graphical user interface and an integrated database management system to aid user analyses. It can accommodate a wide range of support system and aircraft operating characteristics, including the supply and maintenance system echelon structure, the aircraft indenture structure, cannibalization, flying profiles, and support structures.

The ASM is the state of the art in spares requirement models. It is highly capable, yet user-friendly. Its basic principles have been proven in use by the USAF, the National Aeronautical and Space Administration International Space Station Program, and the Israel Air Force. The greater efficiency of the ASM’s optimization can lead to savings of up to 25 percent in spares inventory without degrading support.

Contents

1.0 Introduction	1-1
1.1 The Methodology	1-3
1.2 Benefits of the Weapon-System Approach	1-5
2.0 Basic Weapon-System-Oriented Inventory Systems	2-1
2.1 Overview	2-1
2.2 Base Supply and Item Backorders	2-2
2.2.1 Demand Forecasting	2-6
2.2.2 Base Repair	2-6
2.2.3 The Effects of Depot Stock	2-8
2.2.4 Variance	2-10
2.3 Calculating Availability and Spares Mixes	2-12
2.3.1 Calculating Availability	2-12
2.3.1.1 Quantity per Application	2-13
2.3.1.2 Application Percentage	2-13
2.3.2 Spares Mix Optimization	2-14
3.0 Backorders Under Dynamic Conditions	3-1
3.1 Dynamic Conditions	3-1
3.2 Nonstationary Demand	3-3
3.2.1 The Demand Process	3-3
3.2.2 Backorder Probability Distributions	3-4
3.3 Backorders with Constant Resupply Times	3-6
3.3.1 Introduction	3-6
3.3.2 Assumptions and Definitions	3-6
3.3.3 The Probability Distribution of Base Backorders	3-7
3.3.4 Computational Methods	3-11
3.4 Flying-Hour-Driven Demands and Constant Resupply Times	3-14
3.5 Backorders Under Varying Deterministic Resupply Times	3-16
3.5.1 Resupply Processes with Deterministic Variable Resupply Times	3-16

Contents (Continued)

3.5.2 The Probability Distribution of Base Backorders with Varying Resupply Times	3-17
3.5.3 Flying-Hour-Driven Demands and Variable Resupply Times	3-19
3.6 Other Extensions	3-20
3.6.1 Suspended Resupply	3-20
3.6.2 Levels of Indenture	3-21
3.6.2.1 Introduction	3-21
3.6.2.2 Approach	3-22
3.6.2.3 The Effect of SRU Backorders	3-23
3.6.2.4 Example	3-26
3.6.3 Condemnations	3-27
3.6.4 Days of Warning	3-29
3.6.5 Multiple Bases	3-29
3.6.6 Exponential Repair	3-31
3.6.6.1 Continuous Approach with Stationary Repair Times	3-32
3.6.6.2 Nonstationary Repair Times	3-33
3.6.6.3 Discrete Case	3-35
3.6.6.4 Suspended Resupply and Exponential Repair	3-36
3.6.6.5 Levels of Indenture for Exponential Repair	3-36
3.7 Summary	3-36
4.0 Availability and Optimization	4-1
4.1 Availability with Cannibalization	4-1
4.2 Computing Spares Mixes with Cannibalization	4-3
4.2.1 Confidence Optimization	4-4
4.2.2 The ENMCS Objective Function	4-5
4.2.3 Weighted Heuristics	4-8
4.2.3.1 Weighted Heuristic with an ENMCS Objective	4-9
4.2.3.2 Weighted Heuristic with Joint EBO/ENMCS Objective	4-10

Contents (Continued)

4.2.4	Confidence Optimization with Nonintegral Targets	4-12
4.2.5	Partial Cannibalization	4-13
4.3	Levels of Indenture	4-16
4.3.1	Awaiting Parts Pipelines for LRUs	4-16
4.3.2	Multi-Indenture Optimization	4-18
4.3.3	Common-Component SRUs	4-19
5.0	Sample Results and Sensitivity	5-1
5.1	Model Inputs	5-1
5.1.1	Basic Targets	5-1
5.1.2	Flying-Hour Scenarios	5-2
5.1.2.1	Single-Day Analysis	5-4
5.1.2.2	Two-Day Analysis	5-4
5.1.3	Initial Assets and User-Specified Spares Levels	5-5
5.1.4	Examples of ASM Sensitivity	5-6
5.2	Optimization with Multiple Levels of Indenture	5-6
5.3	The Effects of Cannibalization	5-10
5.3.1	Cannibalization with ENMCS Optimization	5-11
5.3.2	Cannibalization with Confidence or ENMCS/EBO Optimization	5-13
5.3.3	Availability Curves	5-15
5.3.4	Other Model Options	5-16
5.4	Dynamic Flying-Hour Programs	5-17
5.5	Spares Mixes for Peace and War	5-18
5.5.1	Meeting Availability Targets for the Lowest Cost	5-19
5.5.2	Meeting a Fixed Budget Constraint	5-20
6.0	Initial Provisioning	6-1
6.1	Initial Provisioning Time Frame	6-1
6.1.1	Flying Hours in a Coverage Period	6-2
6.1.2	Condemnations in a Coverage Period	6-3
6.2	Common Components	6-4

Contents (Continued)

Bibliography	Biblio. 1
--------------------	-----------

Appendix A. Marginal Analysis

Appendix B. Computing Backorder Statistics: An Example

Appendix C. Recursion Formula for Negative Binomial Distribution

Appendix D. Glossary

Illustrations

1-1.	Availability-vs.-Cost Curve	1-1
1-2.	Basic Model Methodology	1-4
1-3.	How the Weapon-System Approach Affects LRU Spares Quantity	1-6
2-1.	Confidence Level: Poisson Distribution	2-6
2-2.	Cumulative Distribution Function of Demand; Average Demand Equals 4	2-11
2-3.	Generating the Availability Curve	2-17
3-1.	Sample Flying-Hour Profile	3-2
3-2.	LRU/SRU Timeline	3-23
3-3.	SRU Pipeline Computation	3-27
5-1.	Sample Flying-Hour Profile for Aircraft A, B, and C	5-3
5-2.	Differences in Cost and EBOs Between Cannibalizing and Not Cannibalizing for Three NMCS Targets	5-11
5-3.	The Effect of Optimization Methods on the Number of Aircraft Cannibalized in a 20-Aircraft Fleet	5-15
5-4.	ENMCS-vs.-Cost Curves for Three NMCS Targets	5-16
5-5.	ENMCS Over a Period of Dynamic Flying Hours	5-17
5-6.	Peacetime and Wartime Availability — Cost Curves Plotted Separately	5-19
6-1.	Time Frame for Initial Provisioning	6-2
6-2.	Nominal Flying-Hour Program for Initial Provisioning	6-3
6-3.	How Coverage Period Adjusts PLTT	6-4

Tables

1-1.	Item vs. Weapon-System Approach (USAF F-16C Repairable Database)	1-5
1-2.	Typical Kit Savings with ASM Introduction	1-6
3-1.	Peace-to-War Transitional Base Repair Times by Scenario Day	3-20
3-2.	Base Repair Time with Suspension	3-21
3-3.	Peace-to-War Transitional Base Repair Times with Suspension	3-21
4-1.	Comparison of Weights for Confidence and ENMCS Methods (NMCS Target = 4)	4-10
4-2.	Weights for ENMCS and EBO/ENMCS Methods (NMCS Target = 4)	4-11
4-3.	Weights for ENMCS and EBO/ENMCS Methods (NMCS Target = 7)	4-12
4-4.	Weights for Confidence Method (NMCS Target = 3.6)	4-13
5-1.	Resupply Assumptions for F-15E Spares Kit	5-6
5-2.	Spares Mixes for Three Different Indenture Structures (\$2,000 Investment)	5-7
5-3.	Comparing Performance for Three Different Indenture Structures	5-8
5-4.	Comparing LRU/SRU Spares Mix with Respect to Cost	5-8
5-5.	Comparing LRU/SRU Spares Mix with Respect to Pipeline	5-9
5-6.	F-15E Relationship Between LRU and Its SRUs	5-9
5-7.	Reduction in Levels Resulting from Cannibalization (20 Items with Greatest Change)	5-12

Tables (Continued)

5-8.	Changes to Spares Mix Resulting from Cannibalization (\$40 Million Investment)	5-13
5-9.	Results of Different Optimization Techniques (NMCS Target = 4)	5-14
5-10.	Example: When Peace Dominates, with an 80 Percent Availability Input Target for Peace and War	5-20
5-11.	Example: Bounding the Peace and War Availabilities for a \$50 Million Budget	5-20
6-1.	Treatment of Common Component	6-5

1.0 Introduction

Military planners must calculate spares requirements to support weapon-system readiness over a wide range of possible situations. Using the operational parameters of those situations and the characteristics of the weapon system's components — including projected failure rates, repair times, and procurement costs — the Aircraft Sustainability Model (ASM) computes cost-effective spares mixes to minimize waiting time for spare parts. This report documents the technical and mathematical structure of the ASM. It provides a complete description of the ASM: what it does, how it does it, and why it works. It is intended for readers with a reasonable background in probability and statistics. The treatment is otherwise self-contained; familiarity with multi-echelon inventory theory is helpful but not necessary.

In selecting specific spare parts to be procured, the model employs a weapon-system approach, which calls for item decisions to be based explicitly upon their effect on the weapon systems being supported. As the measure of performance, we use the ultimate goal of the logistics system: available aircraft. Availability (at a particular time for a particular aircraft type) is the percentage of the fleet that is not grounded for lack of spare parts. Spares are bought using a marginal analysis technique that purchases items on the basis of their contribution to weapon-system availability per unit cost, thus guaranteeing cost-effective spares mixes. Summary information relating aggregate cost to resulting aircraft fleet availability over the complete range of alternatives is presented to help planners make budgeting and funding decisions (see Figure 1-1). This marginal analysis approach can be applied to any weapon system, from aircraft to air defense radars.




Figure 1-1.
Availability-vs.-Cost Curve

The scope of the model is a single weapon system. An aircraft is “down for supply” or not mission-capable — supply (NMCS) upon failure of a component for which no spare is available. All failures occur at first-echelon sites (bases). The bases are assumed to be uniform with respect to demands, resupply times, and repair capabilities. If an item cannot be repaired at the base, it is shipped to the second echelon (depot) for possible repair, and replenishment from the depot is immediately requested. At the depot, the item may be repaired or condemned. Both echelons — base and depot — are presumed to operate using an $(s - 1, s)$ inventory policy, under which, with every demand, a resupply action is initiated immediately.

The model requires a number of item-specific factors, such as demands (failures) per flying hour, base and depot repair times, the probability of repair at each site, condemnation rates, transportation times, unit cost, quantity per application (QPA), and procurement lead-time.

Also required is the operating tempo (i.e., the weapon system’s flying-hour program). The flying program may be for a steady-state period, for a period with a varying daily flying-hour profile, or for a combination of both. Deriving these flying hours from the more detailed conflict scenarios available to war planners (sortie rates, sortie duration, turn times, etc.) is presumed to be a pre-processing activity.

To produce the relationship between weapon-system availability and funding, this approach considers a number of other factors:

- ◆ *Indenture.* Aircraft repair involves replacing first-indenture components on the aircraft as it sits on the flightline. These first-indenture components are called line replaceable units (LRUs). LRU repair involves replacing second-indenture components in the repair shop; these second-indenture components are called shop replaceable units (SRUs). SRU repair involves replacing third-indenture items, and so on. The model develops the optimal balances between procuring LRU and SRU spares. A spare LRU has a direct impact on availability (when a spare LRU is available, the aircraft is operational almost immediately), while a spare SRU affects availability only indirectly (even if an SRU spare is immediately available to repair the LRU, the aircraft must wait for the LRU to return from maintenance).
- ◆ *Essentiality.* The model focuses on those items that affect availability and cost the most — the “essential” items. These are the high-cost, high-indenture items (e.g., reparable LRUs and reparable SRUs, the items necessary for the weapon system to operate).
- ◆ *Cannibalization.* The ability to consolidate LRU shortages (“holes”) onto a single aircraft greatly improves aircraft availability without increasing procurement costs. For instance, if one aircraft is missing item *A* and another is missing item *B*, maintenance can remove an operational item *B* from the first aircraft and install it on the second, returning that aircraft to service. Cannibalization yields improved aircraft availability over the non-

cannibalization case for the same cost. Similarly, SRUs may be cannibalized during LRU repair. The disadvantage of cannibalization for the maintenance system is that maintenance actions are increased.

- ◆ *Flexible scenarios.* The model can deal with many different operating scenarios and even with changes during the course of the scenario. Operating tempo may be steady state or dynamic; repair and resupply may be suspended for a period; maintenance philosophy (e.g., to cannibalize or not) may change.
- ◆ *Multi-echelon supply structure.* In some cases, the supply system is multi-echelon — inventory is stored at several operating bases and also at a central supply point or depot. Depot inventory is available to all bases but takes a shipping time to arrive. Base stocks are available immediately to satisfy on-site demands but are not intended to satisfy demands at other locations. (Such lateral supply is possible, but typically performed on an ad hoc basis, and it may not be cost-effective. A multi-echelon stockage policy tries to minimize the need for such transfers.) The model allocates stock optimally between the bases and the depot to maximize availability.
- ◆ *Component stock considerations.* The model accepts user-specified spares decisions to handle situations in which assets are already in the inventory or on order from the contractor.

1.1 THE METHODOLOGY

The basic model methodology consists of three steps:

- ◆ The first step involves characterizing the probability distribution of the number of items in various stages of the resupply process (or “pipeline”) — unserviceables in repair at bases or depot and serviceables/unserviceables in transit. The relationship between these quantities and the number and location of spares in the system determines the probability of a backorder.
- ◆ The second step is to relate that item information to weapon-system performance; specifically, to determine the expected number of item backorders, the expected number of aircraft NMCS, and several other weapon-system-oriented measures of supply performance.
- ◆ The third step is to produce the availability-versus-cost curve and the associated optimal spares mix for a specified availability or budget target. The model uses a marginal analysis technique that determines the best mixes of spares for a wide range of targets.

Figure 1-2 displays this three-step process and the possible options within it. In the first step, the model develops item information on the basis of the user’s specification of steady-state or dynamic flying hours. In the second step, the

model uses the item information, from the chosen Step 1 option, to calculate the aircraft availability. (The model uses either of two different methods to calculate availability, depending upon whether or not supply policy permits the possibility of cannibalization of LRUs between aircraft.) The third step then uses the appropriate marginal analysis technique (depending upon the cannibalization policy) to calculate the optimal spares mix.




Figure 1-2.
Basic Model Methodology

Chapter 2 of this report is an overview of the entire three-step process for steady-state conditions without cannibalization. Chapter 2's purpose is to present the entire process simply before moving to the more complex extensions. Chapter 3 extends the presentation to describe the method for computing item information (e.g., the probability distribution of the number of units being repaired or shipped) for a dynamic flying program. Chapter 4 presents the method for estimating availability and performing the optimization with cannibalization for both dynamic and steady-state flying programs. Chapter 5 moves from algorithms to examples and describes how model results vary depending upon input options. Finally, Chapter 6 describes a version of the model to perform a special case of spares procurement — initial provisioning.

1.2 BENEFITS OF THE WEAPON-SYSTEM APPROACH

The advantages of the weapon-system approach are dramatic when compared to the traditional item approach, which typically uses individual item measures to determine how many spares of the item were required. A representative example of this approach is sparing to a "probability of sufficiency", i.e., procure enough spares of the item so that the probability of demand exceeding supply is at least some specified value, say 95 percent. We compared the two approaches using F-16C aircraft data from the United States Air Force (USAF). The results show (see Table 1-1) that, for the same costs incurred under the item approach, the weapon-system approach increased aircraft availability by 30 percent. Alternatively, for the same availability, using the weapon-system approach provided a 40 percent budget savings over the item approach.

Table 1-1.
Item vs. Weapon-System Approach
(USAF F-16C Aircraft Repairable Database)

Performance measure	Item approach	Weapon-system approach	
		Minimizing cost	Maximizing availability
Availability	54%	54%	84%
Cost (\$ millions)	\$14.5	\$8.6	\$14.5

Although the improvements in aircraft availability yielded by using the weapon-system approach are significant, it is interesting to note that, on an item-by-item basis, the actual change in spares levels is usually small. Figure 1-3 illustrates this point. For a given budget, it compares the difference — in spares for LRUs — between the two approaches. Specifically, it subtracts the spares level generated by using the item approach from the level generated by using the weapon-system approach to produce the spares difference. Note that, from the savings generated by reducing a dozen components by one each (those with the negative change shown in Figure 1-3), spares levels for over 300 LRUs were increased between 1 and 6 units over the level generated by using the item approach. (Those dozen LRUs, incidentally, were among the most expensive items but were not the 12 most expensive items.)




Figure 1-3.
How the Weapon-System Approach Affects LRU Spares Quantity

We also estimated the benefits of implementing the ASM for the USAF. Before adopting the ASM, the USAF developed war reserve kits based on a simplified system approach. We estimated that some \$200 million were saved when the USAF adopted the ASM in 1989. Table 1-2 shows the savings for some representative squadrons.

Table 1-2.
Typical Kit Savings with ASM Introduction

Aircraft type (24-aircraft squadron)	Cost (\$ millions)	Aircraft available on day 7
F-16C		
Old kit	16.4	21.7
ASM kit	13.3	22.4
F-15C		
Old kit	14.2	20.0
ASM kit	12.8	20.4

2.0 Basic Weapon-System-Oriented Inventory Systems

This chapter is an introduction to modern weapon-system-oriented inventory system theory. We introduce the standard terminology (in **bold text**) and describe in a simple form the logic and methods used in the ASM. We start at the simplest possible level, describing the interaction between the inventory system and the items it stocks. In the chapters that follow, we build on this foundation, treating more complex situations, until we have covered many of the standard inventory problems and have discussed how the ASM treats them. We present information in a three-step manner: first developing the basic building blocks — the item inventory measures — then addressing the translation of item information into aircraft availability, and finally discussing marginal analysis optimization, which develops the best spares mix on the basis of a spares benefit-to-cost ratio.

2.1 OVERVIEW

The purpose of inventory is to meet needs as they arise. Typical, everyday, inventory decisions are usually made on the basis of intuition or considerations of convenience — carrying no spare tire in your car is too risky, carrying two is wasteful, carrying one is just right. We do not normally have to resort to computer models to master inventory decisions such as these. But larger, more complex systems with larger budgets entail more difficult choices. This is especially true in the case of the aircraft reparable spares we are considering. These components are expensive, some over \$1 million a unit. Many are specialized and not available in the commercial marketplace. Multi-year procurement lead-times are common. And the penalty for shortage is a multimillion dollar weapon system forced to sit idle. Decisions as to the best inventory in such an environment should be explicitly based on cost-effectiveness considerations and are best made using mathematical, statistical models.

The goal of an inventory system is to sustain the operation of an end item — the machinery supported by the spare parts in the system. We are interested here primarily in military aircraft, but application of these concepts to other types of weapon systems or major end items is straightforward. As aircraft are flown they break, and maintenance crews repairing them order spares from the supply system. If this system is performing poorly, it will not be able to supply the needed parts, and many aircraft will be NMCS — that is, grounded for need of a part.¹ A healthy supply system will allow few NMCS aircraft. The

¹Aircraft may also be grounded for want of maintenance — not mission-capable-maintenance (NMCM). Most of this NMCM category results from routine or scheduled maintenance time or from remove-and-replace time for failed components and is not affected by spares inventory.

true measure of supply performance is the ability of the supply system to produce mission-capable aircraft (or other end items). Thus a measure such as **aircraft availability** (the percentage of aircraft not NMCS)² is a good gauge of supply system performance. Certainly, maximizing aircraft availability is the proper goal of the supply system.

Supply system performance has been traditionally measured from the system's internal viewpoint, often by **fill rate** — the percentage of orders **filled** immediately from stock **on hand (OH)** (as distinguished from those that are **backordered** because the supply system is out of stock). One weakness of using fill rate as a measure is that it does not consider backorder duration. Clearly it may make sense to accept a few limited-duration backorder incidents if doing so will prevent the occurrence of rare but lengthy ones. But fill rate is insensitive to this tradeoff. Nor does fill rate capture information about the complexity of the aircraft being supported. All else being equal, more complex aircraft require a higher component fill rate to reach a given availability than do simpler aircraft. We will see later that availability is defined as a product of probabilities — the probability that the aircraft is not missing its first component, times the probability that the aircraft is not missing its second component, and so on. An aircraft with more components has more factors in the product, and since each probability is less than 1.0, the product will tend to be smaller.

Thus using a fill rate criterion, or a backorder criterion alone, for that matter, leads to a bias in favor of the less complex aircraft types. Using a backorder criterion is more defensible if properly done (e.g., in terms of backorders *per aircraft*). As we will show, availability can be viewed as a measure that goes a step beyond backorders, and there is a close relationship between availability and backorders per aircraft. While a high fill rate can be evidence of a healthy supply system, it is not directly related to that system's impact on the ultimate customer — the aircraft. In the difficult cost-effectiveness choices that military logistics planners must make, the difference between fill rate and aircraft availability is critical.

Our goal is to describe how to model an inventory system in terms of aircraft availability. We will begin with a simple system and progress toward a complex multi-echelon, multi-indenture system. The first step toward computing availability is computing item backorder statistics, which we discuss in Section 2.2.

2.2 BASE SUPPLY AND ITEM BACKORDERS

Consider the supply of a single item at an operating base. From time to time a customer (normally an aircraft mechanic) will arrive at base supply and ask for

²The USAF defines NMCS as aircraft grounded for supply only, and defines a separate category for aircraft grounded for want of *both* maintenance *and* supply — NMCB. Our NMCS includes all aircraft grounded for spares, thus embracing the USAF's NMCS plus NMCB. The USAF refers to our NMCS as TNMCS (total not mission-capable — supply).

a unit of the item. If base supply has the item in stock, it fills the customer's order; if not, the request is backordered and the customer must wait until a unit becomes available.

Suppose that the base is operating in a steady state, typical of peacetime operations. While flying hours and sorties may not be identical from day to day, there are no great changes as would be seen in a wartime or contingency deployment. Then the arrival of demands can be described by a stationary random (stochastic) process. Frequently a Poisson process is used to describe this demand activity, though a negative binomial characterization — which embodies more variability — seems more descriptive of USAF activity. At any rate, for a stationary Poisson process, if the expected number of **demands** per day is represented by λ , the expected number of demands in a time period, T , is λT . For a Poisson process, the probability that *exactly* n demands will occur in T days is

$$p(n) = \frac{e^{-\lambda T} (\lambda T)^n}{n!}. \quad [\text{Eq. 2-1}]$$

In this chapter, "Poisson process" means stationary Poisson process. Chapter 3 treats nonstationary processes.

Suppose there is no base repair capability and that, instead, base supply is supported by a **higher echelon** (a **wholesale** supplier — the **depot**). With each failure, the base ships the unserviceable unit back to the depot while simultaneously ordering a replacement. The **resupply time** — how long the base must wait for the unit to arrive — is critical to the base's ability to support its customers and to the calculation of how many spare units it should stock.

For now we suppose that the depot always has adequate stock, so that any order received from a base is always filled and never backordered. Then the base's resupply time is just the time for the depot to process the order and ship it to the base, the **order and ship time (OST)**. We assume that OST is constant; variable resupply times are considered in Chapter 3.

If base supply orders a unit from the depot every time a demand occurs, there will be a stream of units flowing from the depot to the base, an example of a **pipeline**. This particular pipeline is the **order and ship pipeline (OSpipe)**. Since the number of units in this pipeline at any time, t , is simply the number of demands between time $t - \text{OST}$ and time t , Equation 2-1 tells us that the number of units in the pipeline has a Poisson probability distribution with mean $\lambda \times \text{OST}$. The number of items in the order and ship pipeline is a random variable, **OSpipeRV**. We will denote the mean of this pipeline by **OSpipe**. Because the number of items in the order and ship pipeline is a Poisson random variable, its distribution is completely determined by this mean.

Pipeline probability distributions are important because once we can describe the pipeline's behavior, we can measure how well base supply is doing its job. For example, the fill rate can be characterized as the probability that base supply has a unit on hand at a random point in time. To compute this rate, we

need only know the probability distribution of the number of units in the pipeline and the base's **inventory**, as we will see.

Before we can do this, however, we need to explain what we mean by the base's inventory of an item. Not every unit "owned" by base supply is sitting on the shelf ready for issue. Base supply might own six units, but five might be currently in the order and ship pipeline (leaving only one on the shelf). In fact, if eight are in that pipeline, base supply would have none on the shelf and two backordered demands waiting to be filled.

The total number of units owned by base supply (including any units in resupply pipelines) is called the base's **inventory position**, or spares level, s . The inventory position is defined as the number on hand (OH), plus the number **due in** (DI) (i.e., in a pipeline), minus the number **due out** (DO) (i.e., backorders owed to customers). Thus

$$s = OH + DI - DO. \quad [\text{Eq. 2-2}]$$

Often the term **backorder** (BO) is used instead of **due out** (DO).

We will assume that the base operates under an $(s-1, s)$ **inventory policy**, which dictates that when a demand lowers the inventory from s to $s-1$ (which is the **reorder point**), it triggers an action that immediately raises the inventory position back to s . Less expensive consumable items — filters, belts, nuts and bolts, etc. — are often managed under the more general (r, Q) **inventory policy**. Under an (r, Q) inventory policy, the inventory position is allowed to fall until it hits r (the reorder point), at which time the **order quantity**, Q , is ordered from the wholesaler to bring the inventory position up. With an $(s-1, s)$ inventory policy, the inventory position is *always* equal to s , and we refer to the inventory position or spares level as the constant s .

Under our $(s-1, s)$ assumption, we see that in the previous example, the number in the pipeline (DI), the number on the shelf (OH), and the number of backorders (DO) can vary, but the inventory position stays fixed at 6. For example, if 5 are DI and the base gets a demand, it is filled with the last OH unit, dropping the base's inventory position to 5 ($0 + 5 - 0$). But base supply immediately orders another unit from the depot, bringing the inventory position back to 6 ($0 + 6 - 0$). If another demand arrives, it cannot be filled, causing a backorder (DO), again lowering the inventory position to 5 ($0 + 6 - 1$) — but once again, an order to wholesale immediately restores the inventory position to 6 ($0 + 7 - 1$).

Nor does the inventory position change when an order arrives from wholesale (i.e., emerges from the order and ship pipeline). We left the example above with 1 due out, and $s = 0 + 7 - 1 = 6$. When an order arrives, the number in the pipeline changes from 7 to 6, but the number on hand rises to 1, so $s = 1 + 6 - 1 = 6$. Base supply immediately fills the backordered demand, and $s = 0 + 6 - 0 = 6$ still.

Knowing the inventory position and the probability distribution of the number of units, x , in the pipeline, we can compute several widely used measures of supply performance. When x is less than s , there are units on hand ($OH = s - x$). When x is greater than s , there are backorders ($BO = x - s$). Thus the fill rate is just the probability that $x < s$:

$$\text{fill rate} = \sum_{x=0}^{s-1} p(x), \quad [\text{Eq. 2-3}]$$

where $p(x)$ is the probability of x units in the pipeline.

However, as we explained earlier, fill rate is only a mediocre measure of military supply effectiveness. A better measure — one that moves us closer to computing aircraft availability — is the mean or expected number of backorders (**EBO**). Since $BO = x - s$ for x greater than s ,

$$EBO = \sum_{x>s} (x - s)p(x). \quad [\text{Eq. 2-4}]$$

Another commonly used measure is the probability of no backorders (also known as the **ready rate** or **confidence level**) as an indication of supply effectiveness. Confidence level is given by

$$\text{confidence level} = \sum_{x=0} p(x). \quad [\text{Eq. 2-5}]$$

The confidence level for s spares equals the probability of s or fewer units in resupply, which is just the **cumulative distribution function** $CDF(s)$. Figure 2-1 shows the confidence level as a function of the number of spares for a component whose expected pipeline equals four. The chart shows the **probability distribution function**, $PDF(x)$, and the cumulative distribution function, $CDF(x)$, of the random variable x , the number of units in resupply. [Note: $CDF(x) = CDF(x - 1) + PDF(x)$.] In Figure 2-1, a confidence level of 95 percent requires seven spares.




Figure 2-1
Confidence Level: Poisson Distribution

2.2.1 Demand Forecasting

The mean daily demand, λ , is a critical model input. In accordance with the USAF policy in effect when the ASM was developed, λ is estimated as the product of the failure factor (FF)³ — the historical demands per flying hour — and the total daily flying-hour program (FHP). Recent LMI research has corroborated the widely held belief that flying hours are not a very good indicator of demands. Using sorties per day as an indicator is typically better overall, though flying hours do have some predictive value [Sherbrooke, 1996].⁴ Other program units are more appropriate for other classes of components, (e.g., projected rounds fired for gun parts).

Judicious preprocessing can be used to replace flying hours with more appropriate programs to obtain the best estimate of mean daily demand.

2.2.2 Base Repair

Suppose that the base, instead of always having to order another unit from the depot, can resupply itself by fixing the broken part turned in by maintenance. The broken carcass is **inducted** into **base repair**, where it is either fixed or declared **not reparable this station (NRTS)**. If the part is NRTS, the base orders a unit from the depot and sends the carcass back to the depot for repair.

³ Failure factor is sometimes referred to as the total organizational and intermediate demand rate (TOIMDR) in Air Force supply information systems.

⁴ Sources are listed alphabetically by author's last name in the bibliography.

The base now has a **base repair pipeline (BRpipe)** in addition to the order and ship pipeline. We denote the number of units in the pipeline, a random variable, by $BRpipeRV$, and the pipeline mean by $BRpipe$. If we know the **NRTS rate** (the percentage of demands that are NRTS, resulting in demands on the depot) and the **base repair time (BRT)**, then we can calculate the probability distributions for the number of items in the base repair and order and ship pipelines and incorporate them into the backorder calculations. As with OST, we assume that BRT is constant.

We may think of each demand on the base as having a probability $NRTS$ of resulting in a demand on the depot and a probability $1 - NRTS$ of resulting in a demand on base repair. Since we are assuming that the original demand process is Poisson, we can use the splitting theorem for such processes [Feller, Vol. 1] to conclude that demands on base repair and demands on the depot are independent Poisson processes with mean daily demand rates $(1 - NRTS) \times \lambda$ and $NRTS \times \lambda$, respectively. Using Equation 2-1, as we did for the order and ship pipeline earlier, we see that the number of items in the base repair pipeline and the number of items in the order and ship pipeline are Poisson random variables with means given by Equations 2-6 and 2-7, respectively.

$$BRpipe = \lambda \times (1 - NRTS) \times BRT. \quad [Eq. 2-6]$$

$$OSpipe = \lambda \times NRTS \times OST. \quad [Eq. 2-7]$$

Since the demand processes on base repair and on the depot are independent [Parzen], it follows that the number of items in the base repair pipeline and the number of items in the order and ship pipeline are independent random variables.

Because the numbers of items in the two pipelines are independent Poisson random variables, the number of items in the total base resupply pipeline is a Poisson random variable with mean equal to sum of the means of the two pipelines (this follows from convolving the two pipeline distributions). Thus the mean number of items in the total base resupply pipeline is

$$Bpipe = \lambda \times [NRTS \times OST + (1 - NRTS) \times BRT]. \quad [Eq. 2-8]$$

Since the total number of units in the base resupply pipeline has a Poisson probability distribution, that distribution is completely determined by its mean, $Bpipe$. Therefore we can compute the probability $p(x)$ that there are x units in the total base resupply pipeline for any x , and, by substituting those probabilities into Equation 2-4, we can compute the expected backorders at the base as a function of the number of spares in the base stock level s . Knowing each item's expected backorders as a function of its stock level is the first step in computing aircraft availability as a function of stock level.

2.2.3 The Effects of Depot Stock

So far we have assumed that the depot could always satisfy demands with stock on hand; in effect, we assumed unlimited serviceable depot stock. We now assume that depot stock is finite, that the depot may repair items, and that each item's depot repair time is constant. Here, "depot repair time" includes a retrograde time to ship an item from the base to the depot. This section examines the effect of these new assumptions on the total base resupply pipeline. We continue to focus on a single item.

Depot stock fluctuates as does base stock except that the depot's customers are bases instead of mechanics. When a base orders a part, either the order is filled or the item is backordered, depending on whether the depot has a serviceable unit on hand. The base turns in a broken carcass, which goes into **depot repair**. There is now a **depot repair pipeline** as well as a base repair pipeline. We denote the number of units in that pipeline by DRpipeRV and the pipeline mean by DRpipe. We assume for now that no items are condemned (i.e., that the system is **conservative**).

Assume now that we have N bases and that the item's NRTS rate and order and ship time are uniform across bases. We assume that the depot follows an $(s - 1, s)$ inventory policy, as do the bases. Suppose the demand on base b is a Poisson process with mean λ_b , and the demands on the bases are mutually independent.

Then base b 's demand on the depot is a Poisson process with mean $\lambda_b \times NRTS$, as we saw in section 2.2.2, and the demands of the bases on the depot are mutually independent. Therefore the total demand on the depot is Poisson with mean $\lambda_0 = \sum_{b=1}^N \lambda_b \times NRTS$. Letting DRT be the constant depot repair time, and using Equation 2-1, we see that DRpipeRV is Poisson-distributed with mean $\lambda_0 DRT$. We can now compute the depot expected backorders (DEBO) using Equation 2-4, interpreting s as the depot stock level, x as the number of items in the depot repair pipeline, and $p(x)$ as the probability that there are x units in the depot repair pipeline. We will use DBORV to denote the number of depot backorders, a random variable not usually Poisson distributed.

Now consider a particular base, base j . We will compute the probability distribution of the number of items in base j 's resupply pipeline; we denote the number in that pipeline by BpipeRV _{j} and its mean by Bpipe _{j} . Let BRpipeRV _{j} and OSpipeRV _{j} be the number of units in base j 's base repair and order and ship pipelines, respectively; let BRpipe _{j} and OSpipe _{j} be the means of these pipelines. Let DBORV _{j} be the number of backorders at the depot representing due ins to base j and DEBO _{j} be the mean number of these due ins.

Since a unit due in to base j is either in base repair, in transit from the depot, or backordered at the depot, we see that the number of units in base j 's resupply pipeline is just

$$BpipeRV_j = BRpipeRV_j + OSpiceRV_j + DBORV_j. \quad [\text{Eq. 2-9}]$$

Because the mean of a sum of random variables is just the sum of their means, the mean number of units in the total base resupply pipeline is

$$Bpipe_j = BRpipe_j + OSpice_j + DEBO_j. \quad [\text{Eq. 2-10}]$$

We assume that the mean number of backorders at the depot owed to each base is in proportion to that base's fraction of total demand;⁵ that is,

$$DEBO_j = \frac{\lambda_j}{\lambda_0} DEBO. \quad [\text{Eq. 2-11}]$$

In the current version of the model, we assume that $\lambda_j = \lambda$ for all j , so that we have

$$Bpipe_j = BRpipe_j + OSpice_j + \frac{DEBO}{N}. \quad [\text{Eq. 2-12}]$$

At this point, we use one of several approximations (the distribution of $BpipeRV_j$ is not usually Poisson) to compute the probability distribution for the base resupply pipeline for a single base, depending on the user's choice of the variance-to-mean ratio option. The default ASM option uses a constant variance-to-mean ratio of 1 to approximate the probability distribution of each base's resupply pipeline by a Poisson distribution with mean $Bpipe_j$; for other options, we approximate the probability distribution using the variance of the single base resupply pipeline as well as its mean. We discuss the effects of these options in Section 2.2.4 and in Chapter 3.

We next substitute the probability distribution of the base resupply pipeline for $p(x)$ in Equation 2-4 to compute the single base expected backorders as a function of both the base and depot stock levels (recall that $DEBO$ in Equation 2-12, and thus $p(x)$ in Equation 2-4, depends on the depot stock level).

For each number of spares, the model finds the allocation of that number of spares across the depot and the N bases that gives the lowest number of base expected backorders; it does this by exhaustively trying all of the combinations of depot and base spares that sum to the given number of spares. The expected backorders for the allocation of each number of spares across the depot and the bases is stored and later used in the availability optimization described in Section 2.3.2.

⁵See Chapter 3 for further discussion.

The ASM can also allow for condemnations (i.e., some carcasses may not be reparable). For most expensive reparables, the fraction of demands that are condemned, **ConPCT**, is small and in many contexts may be ignored. For a steady-state formulation, we may assume an $(s - 1, s)$ procurement policy to replace condemnations, though the USAF actually uses a periodic review system, typically initiating procurement actions only once a year. In the $(s - 1, s)$ formulation, the depot pipeline gains another segment, and the mean number of units in the depot resupply pipeline is then expressed as

$$DR_{pipe} = \lambda_0 \times [ConPCT \times PLT + (NRTS - ConPCT) \times DRT], \quad [Eq. 2-13]$$

where **PLT** is the **procurement lead-time**. Using this new mean for the depot resupply pipeline, we compute the probability distribution for the number of units in that pipeline and expected depot backorders as before; the computation of expected backorders at a single base is unchanged as well.

2.2.4 Variance

Thus far, we have considered Poisson processes only. Actual demands can be more erratic than a Poisson process. We may also want to consider **uncertainty** about the demand rates caused by the inherent difficulty in forecasting. In either case, we must use a probability distribution that has more **variance** than a Poisson distribution. The practical choice for computer-based models is the **negative binomial distribution**. While the Poisson is a single-parameter distribution whose variance always equals its mean, the negative binomial distribution may have a variance larger than the mean. With a negative binomial distribution, Equation 2-1 for the probability of n demands in a time interval of length T is replaced by

$$p(n) = \frac{Q^{\frac{\lambda T}{P}} \left(-\frac{P}{Q}\right)^n \Gamma\left(n - \frac{\lambda T}{P}\right)}{n! \Gamma\left(-\frac{\lambda T}{P}\right)}, \quad [Eq. 2-14]$$

where λT is the mean, Q is the variance-to-mean ratio (VMR), $P = 1 - Q$, and Γ is the gamma function.

In the case of a constant $VMR > 1$, the ASM uses Equation 2-14 to describe an item's demand distribution in a time interval T . The number of units in each pipeline segment, except for depot backorders,⁶ is described by negative binomial distributions with means as calculated previously and specified VMR. The number of units in each pipeline segment are treated as independent random

⁶Even if the underlying demand process is Poisson, the number of depot backorders is not Poisson distributed unless the depot spares level is zero. In that case, the number of depot backorders equals the number of units in the depot repair pipeline. If the depot spares level is $s > 0$, then the probability distribution of depot backorders is equal to the tail of the (Poisson) probability distribution of the depot repair pipeline, translated by s , since the probability of x depot backorders is the probability of $x + s$ units in depot repair. We discuss this further in Chapter 3.

variables, so that the mean and variance of the total base resupply pipeline are the respective sums of the means and variances of the pipeline segments. We then use this mean and variance as parameters of a negative binomial distribution and use that distribution to approximate the probability distribution of the number of units in the total base resupply pipeline.

The ASM may optionally use the approach of Vari-METRIC [Slay]. In this case, the demand process is treated as Poisson but with an uncertain mean with a gamma distribution. This leads to a negative binomial distribution for demands in a time interval. The individual pipeline segments are not independent, since the gamma prior distribution for the demand on the base affects each of the pipelines. The Vari-METRIC option explicitly considers pipeline correlation in computing the pipeline variances and models the effect of depot spares on the variance of the total base resupply pipeline. We discuss this option further in Chapter 3.

The fact that the negative binomial provides a two-parameter family of distributions allows the approximation of a wide range of distributions (by estimating their means and variances).

Figure 2-2 shows various demand distributions⁷ with a mean of 4, but VMRs of 0.6, 1, and 3. The number of spares required to obtain a 95 percent confidence level equals 6, 7, and 10, respectively. In other words, when all else is equal, more spares are required by items with greater demand uncertainty (larger VMRs). The VMR is usually set on the basis of a forecasting standard or in accordance with an empirically derived formula.




Figure 2-2
Cumulative Distribution Function of Demand; Average Demand Equals 4

⁷We have discussed the Poisson distribution (VMR=1) and the negative binomial distribution (VMR>1); a demand distribution with a VMR less than one can be modeled with a binomial distribution. The current version of the model does not handle this case, which rarely occurs.

2.3 CALCULATING AVAILABILITY AND SPARES MIXES

In this section, we start with the item-oriented measures of pipeline and backorders (expected value and probability distribution) just developed and then, by considering their impact on the aircraft, develop our performance measure — aircraft availability. Once we have a performance measure, we can then rank each spare on the basis of a benefit-to-cost ratio and develop an optimal spares mix.

2.3.1 Calculating Availability

In this section, we discuss how to use the item backorder distribution we have derived to estimate aircraft availability. First we discuss the estimation of availability assuming no cannibalization. Discussion of availability when there is cannibalization is discussed in Chapters 3 and 4.

We assume for now that all the components of the aircraft are LRUs. Thus they are removed directly from the aircraft upon failure, and a spare is needed to return the aircraft to operational status. SRUs must also be considered because their shortages may delay LRU repairs, extend effective resupply times, and increase LRU backorders. Rather than complicating this stage of the exposition, however, we defer discussion of SRUs to Chapters 3 and 4.

If a part on an aircraft breaks and supply is out of stock for that part, then the part is backordered. This backorder causes a hole on the aircraft, and the plane is NMCS. Since aircraft availability is the percentage of planes not NMCS, computing aircraft availability involves calculating the number of backorders (or holes on the aircraft) and how many planes those backorders make NMCS (how many different planes have holes).

If each aircraft had only one part on it, the formula for aircraft availability would be simple. Each backorder grounds one aircraft, so the number of NMCS aircraft equals the number of backorders. If we observe a number of backorders at a specific time, then observed aircraft availability is given by

$$\text{observed availability} = 1 - \frac{\text{NMCS}}{\text{NAC}} = 1 - \frac{\# \text{ of backorders}}{\text{NAC}}, \quad [\text{Eq. 2-15}]$$

where NAC is the number of aircraft in the fleet. Thus if there were 2 backorders and 10 aircraft in the fleet, the availability would be 0.8, or 80 percent.

The expected availability, in which we are usually interested, is obtained by applying the expectation operator in Equation 2-15, yielding

$$\text{availability} = 1 - \frac{\text{EBO}}{\text{NAC}} = 1 - \frac{E(\text{NMCS})}{\text{NAC}}. \quad [\text{Eq. 2-16}]$$

We can interpret availability as the probability that a randomly selected aircraft is not NMCS.

The extension to multiple components is straightforward if we assume independence of the backorder process across components. Let EBO_l be the expected backorder total for LRU l , with a given spares level. Because backorders of different types of components are independent, the probability that a randomly selected aircraft is not NMCS for any LRU l , is the product, over all the LRUs, of the probabilities that the aircraft is not NMCS for each one. Thus

$$availability = \prod_l \left(1 - \frac{EBO_l}{NAC} \right). \quad [Eq. 2-17]$$

2.3.1.1 QUANTITY PER APPLICATION

Some components are installed in the aircraft in pairs, triples, etc. For example, many aircraft have more than one engine. To account for this, we let QPA_l be the quantity per application and let TI_l be the **total installed** for part l . Then ($TI_l = NAC \times QPA_l$) is the number of possible locations for a hole. Note that TI_l plays the role that the number of aircraft played in Equations 2-15 through 2-17, where we could have at most one hole for a given component per aircraft. Thus the expected availability is

$$availability = \prod_l \left(1 - \frac{EBO_l}{TI_l} \right)^{QPA_l}. \quad [Eq. 2-18]$$

We could have modeled availability with a more complicated formula that took into account the fact that the first backorder for item l could occur in any of TI_l slots, the second in any of $TI_l - 1$ slots, etc.⁸ However, since the availability targets we typically use require LRU backorders to be small, the chance of a given slot being assigned more than one backorder can be ignored. Moreover, correcting for this possibility would add tedious combinatorial arithmetic to the model.

2.3.1.2 APPLICATION PERCENTAGE

Now consider the case where a component is applied to only part of an aircraft fleet [the application percentage (AP),⁹ expressed as a fraction, is less than 1.0]. Given that an aircraft does not have component l , the probability that it is not NMCS for component l is 1.0 — such an aircraft cannot be grounded for lack of that component. The probability that an aircraft randomly selected from those with component l is not NMCS for component l is $\left(1 - \frac{EBO_l}{TI_l} \right)^{QPA_l}$, where $TI_l = NAC \times QPA_l \times AP_l$. The probability that an aircraft chosen at random has component l is AP_l and the probability that it does not have component l is $1 - AP_l$.

⁸For a more precise treatment, see Gaver.

⁹In typical applications, we project forward to a point where the spares we decided to buy have been delivered, and characterize aircraft and configuration then. Thus the term future application percentage (FAP) is often used.

We assume that the presence or absence of components l and k are independent for $l \neq k$. While this is not always true in practice, we have found this a reasonable approximation. Let A_l be the probability that a randomly selected aircraft is not NMCS for component l ; let $X = 1$ if that aircraft is not NMCS for component l and $X = 0$ otherwise; let $Y = 1$ if that aircraft has component l and $Y = 0$ if not. Then the probability that a randomly chosen aircraft is not NMCS for component l , A_l , is

$$\begin{aligned} \Pr(X = 1) &= \Pr(X = 1|Y = 0)\Pr(Y = 0) + \Pr(X = 1|Y = 1)\Pr(Y = 1) \\ &= 1 \times (1 - AP_l) + \left(1 - \frac{EBO_l}{TI_l}\right)^{QPA_l} \times AP_l. \end{aligned} \quad [\text{Eq. 2-19}]$$

From the independence of backorders for different components and of the presence of distinct components, the availability of a randomly chosen aircraft based on all components is

$$availability = \prod_l A_l = \prod_l \left[(1 - AP_l) + AP_l \left(1 - \frac{EBO_l}{TI_l}\right)^{QPA_l} \right]. \quad [\text{Eq. 2-20}]$$

Note that considering these more detailed data elements requires straightforward changes to several computations. In particular, $TI_l = NAC \times QPA_l \times AP_l$ and component l 's demand rate $\lambda_l = FHP \times QPA_l \times AP_l \times FF_l$, where FHP is the total fleet flying hours per day, and FF_l , the failure factor, is the expected number of failures per flying hour for component l .

2.3.2 Spares Mix Optimization

Given spares levels for each component on an aircraft, we now can compute the expected availability for the fleet and the contribution of each component to that availability. The final step is to determine the best spares mix. For that step we employ a marginal analysis technique that ranks additional spares of each item on the basis of their contribution to weapon-system availability per unit cost. Buying from this ranked list in order, until funds are exhausted, attains the maximum availability possible for a given resource level.

Suppose we have an initial spares level N_l for each LRU $l = 1, 2, \dots, L$ and a corresponding total $EBO(l, N_l)$ that is the minimum expected base backorders achieved by an "optimal" allocation of the N_l spares among the depot and bases. Then the expected availability for the fleet is given by

$$A = \prod_l A_l(N_l), \quad [\text{Eq. 2-21}]$$

where $A_l(N_l)$ is the expected item availability for component l with N_l spares as in Equation 2-20 — $A_l(N_l)$ is the probability that a randomly selected aircraft is not missing component l .

Suppose we add another unit of component j to our inventory. Then the expected number of backorders for component j is reduced from $EBO(j, N_j)$ to $EBO(j, N_j + 1)$, and the new availability is given by

$$\begin{aligned} A' &= \left[\prod_{l \neq j} A_l(N_l) \right] \times A_j(N_j + 1) \\ &= A \times \frac{A_j(N_j + 1)}{A_j(N_j)}. \end{aligned} \quad [\text{Eq. 2-22}]$$

Thus

$$\frac{A'}{A} = \frac{A_j(N_j + 1)}{A_j(N_j)}. \quad [\text{Eq. 2-23}]$$

Similarly, in the general case, with any existing spares mix $\{n_i\}$ and resulting availability A , increasing the spares level of component j from n_j to $n_j + 1$ yields a new availability A' , where

$$\frac{A'}{A} = \frac{A_j(n_j + 1)}{A_j(n_j)}. \quad [\text{Eq. 2-24}]$$

Note that the improvement in availability depends only on component j — the availability function is separable.

Thus we can define, for every component l and every integer n_l greater than the initial spares level, an "improvement factor" — the increase in availability that results from increasing the level of component l from $n_l - 1$ to n_l ; namely, $I(n_l) = \frac{A_l(n_l)}{A_l(n_l - 1)}$. We can thus form an array for each component l , consisting of $EBO(l, n)$, $A_l(n)$, and $I(n)$ for $n = n_l + 1, n_l + 2, \dots$ until $EBO(l, n)$ is effectively zero. We can add to this array the marginal benefit per dollar of adding each additional spare. We call this the sort value of the n th spare of component l , $V(l, n)$, and its exact form is given by

$$V(l, n) = \frac{\ln[I(l, n)]}{C_l}, \quad [\text{Eq. 2-25}]$$

$$V(l, n) = \frac{\ln A_l(n) - \ln A_l(n - 1)}{C_l}. \quad [\text{Eq. 2-26}]$$

The introduction of the natural logarithm may be surprising. In Appendix A, we discuss in detail the optimization of separable functions with marginal analysis and apply that technique in several contexts occurring in the ASM. In the typical marginal analysis approach, the separable multi-variable objective function is expressed as a sum of single variable functions. In this case, the objective function is additively separable, as opposed to the multiplicatively separable availability function in Equation 2-22. The use of the natural logarithm converts the availability function into a form suitable for the marginal analysis machinery.

We can then sort all of the spares that are candidates for procurement in descending order of $V(l, n)$. With N_i representing the initial asset level of each component, the starting availability with no procurement is

$$A = \prod_i A_i(N_i). \quad [\text{Eq. 2-27}]$$

If j is the component with the highest sort value, the first procurement action should be to increase the spares level of component j from N_j to $N_j + 1$. The resulting availability, A' , can be calculated as

$$A' = A \times \exp[C_j \times V(j, N_j + 1)] \quad [\text{Eq. 2-28}]$$

as we see below:

$$\begin{aligned} A \times \exp[C_j \times V(j, N_j)] &= A \times \exp\left\{C_j \times \frac{\ln[I(j, N_j + 1)]}{C_j}\right\} \\ &= A \times I(j, N_j + 1) \\ &= A'. \end{aligned} \quad [\text{Eq. 2-29}]$$

Similarly, as we add spares, going from, say, a level of $n - 1$ to n for component j , the new availability, A' , can be expressed in terms of the old availability, A , the sort value, and the item cost by

$$A' = A \times \exp[C_j \times V(j, n)]. \quad [\text{Eq. 2-30}]$$

As we add spares, we track accumulating costs and the resulting fleet availability. The list of spares added is the **shopping list**. The availabilities and costs arising from each added spare produce a curve, as shown in Figure 2-3. The mathematics of marginal analysis described in Appendix A ensures that buying from this list until our funds are exhausted provides the maximum availability. Alternatively, if we have a desired target availability, we can buy down the list until that availability is reached and know that we have a minimum-cost solution. The curve provides macro-level summary information to planners for budget formulation and allocation.




Figure 2-3.
Generating the Availability Curve

The next chapter examines the expected backorder computation under dynamic conditions; in Chapter 4, we further discuss optimizing the spares mix, allowing for cannibalization.

3.0 Backorders Under Dynamic Conditions

This chapter has several sections. Section 3.1 discusses what we mean by “dynamic conditions,” tells why we need to consider them, and contrasts them with those assumed by the steady-state model in Chapter 2. The steady-state backorder computation described in that chapter should be regarded as a special case of the more general case presented here (and one easier to understand); it is not necessary to employ two distinct models for the steady-state and dynamic cases. Sections 3.2 and 3.3 explain the assumptions and mathematics we use to handle the dynamic case and the reasons why the estimation of backorder statistics is the goal of this chapter. In Section 3.4 we introduce a simple dynamic case by allowing the flying program (which we assume drives the demand process) to vary, but we continue to assume constant resupply times. Further sections relax the assumption of constant resupply times, and discuss other generalizations, such as suspended resupply and SRU backorders.

3.1 DYNAMIC CONDITIONS

In Chapter 2, we considered **steady-state** conditions; that is, we supposed that flying hours or sorties, which we assume to be the drivers of an item’s demand process, were constant over time and that the parameters characterizing the inventory systems’ response to demands (OSTs, repair times, etc.) were constant as well. Our method of computing an item’s expected backorders, the crucial first step in estimating aircraft availability, depended on the steady-state assumption.

In this chapter, the steady-state backorder calculations described in Chapter 2 will be generalized in order to handle **dynamic conditions** — that is, situations in which the flying program, and hence the demand process, changes over time. (With the steady-state assumption, demand varied from day to day, being random, but the underlying process generating that demand did not.) We allow the inventory system parameters characterizing the response of that system to demands (e.g., OSTs, repair times, and NRTS rates) to change as well. These dynamic conditions require **dynamic backorder** calculations that explicitly consider changes in the demand generation process and in the time lags that occur as spares move through the various resupply pipelines.¹ In Chapter 4, we discuss the computation of availability and an economically efficient spares mix under more general conditions than those of Chapter 2.

One situation that leads us to consider dynamic conditions is planning for wartime, when the use of aircraft (i.e., the flying hours, sorties, or other

¹For another treatment, see Isaacson et al.

measures of activity) changes rapidly from day to day, and thus so does the component demand process. In the transition from peace to war, resupply times may shorten or lengthen; resupply may even be suspended. Dynamic conditions may also accompany the introduction of a new aircraft or a force drawdown following a war (increasing or decreasing flying programs, respectively).

A typical application of the model — a useful one for the USAF — assumes a period of steady-state conditions followed by a period of dynamic conditions. That scenario portrays the less strenuous flying hours of peacetime followed by the surge in activity of a conflict. Figure 3-1 displays a typical flying-hour profile for such a case. The profile represents the total number of flying hours per day for the entire aircraft squadron. Our time period convention is that the steady-state period goes from a “negative” day (far enough back to reach steady-state conditions as we have described them) to day zero, while the dynamic period starts on day 1 and ends as much as 99 days later. All days preceding day zero have the same (peacetime) flying hours as day zero, but flying hours for days after day zero may vary.




Figure 3-1.
Sample Flying-Hour Profile

Nothing in our approach requires that the program driving demand be flying hours, although that has been the program used in the USAF case; the program driving demand could be sorties, gun firings, or some other program that has empirically been shown to have a predictive relationship with the average demand rate.

Under dynamic scenarios, the performance measure is not usually taken to be aircraft availability on a single day, as there may be several crucial days on which we need to measure supply performance. Furthermore, because of

changes in repair and resupply, the optimal mix of spares (see Chapters 2 and 4) may vary with the day on which we seek to maximize aircraft availability. Therefore our measure of supply performance is *sustainability*, which typically consists of the aircraft availability measured on two or more days.

Whatever scenario we are considering, our notion of sustainability still requires aircraft availability according to day, and the first step toward computing aircraft availability is to obtain the *time-dependent probability distribution of backorders*.

3.2 NONSTATIONARY DEMAND

In this section, we briefly review the steady-state mathematical assumptions of Chapter 2 and their significance for the aircraft availability computation described there. We then contrast those assumptions with the ones we use to compute availability under dynamic conditions.

3.2.1 The Demand Process

In Chapter 2 we assumed that an item's demand on base supply could be reasonably approximated by a stationary (strict-sense stationary) Poisson process; that is, we assumed that the probability of k demands in *any* time period of length T was given by $\frac{(\lambda T)^k}{k!} e^{-\lambda T}$, where λ was the long-term average demand rate. The steady-state assumption was reflected in the fact that the probability of any given number of demands in a time interval depended only on the length of the interval. The assumption of a Poisson probability distribution for the number of demands in an interval resulted in less variance than is seen in most USAF demand data. However, it allowed us to use the simple result about the splitting of a stationary Poisson process [Feller, Vol. 1] to conclude that the base resupply pipelines were independent Poisson processes in their own right.

In this chapter, we assume that the demand on base supply can be approximated by a nonstationary Poisson process (the probability of a given number of demands in a time interval depends on both the length and on the location of the time interval).

In detail, let $N(t)$ = the number of demands in the time interval $[0, t]$. We say that $N(t)$ is a **nonstationary** (or **nonhomogeneous**) **Poisson process** if the following conditions hold:

- ◆ $N(0) = 0$.
- ◆ If I_1 and I_2 are non-overlapping time intervals, and N_1 and N_2 , respectively, are the number of demands occurring in those time intervals, then N_1 and N_2

are independent random variables (a process satisfying this condition is said to have **independent increments**).

- ◆ There is a function $\lambda(t)$ such that the probability distribution for the number of demands in any time interval (t_1, t_2) is Poisson with parameter $\Lambda(t_2) - \Lambda(t_1)$, where $\Lambda(t) = \int_0^t \lambda(s)ds$.

The function $\lambda(t)$ is called the *intensity* of the process and may be interpreted as follows: the probability of a single demand in a short time interval $(t, t + \Delta t)$ is $\lambda(t)\Delta t + \varepsilon(t)$, where $\frac{\varepsilon(t)}{\Delta t} \rightarrow 0$ as $\Delta t \rightarrow 0$. For a physical interpretation, think of intensity at a point in time as measuring the likelihood of a demand occurring in a short interval containing that time; the larger the intensity is, the greater is the likelihood of a demand. The function $\Lambda(t)$ is known as the *mean value* function of the process and may be thought of as the mean of the number of demands in the interval $(0, t)$. It is straightforward to show that if $\lambda(t) = \lambda$ for all t , the process is simply the stationary Poisson process of Chapter 2.

The assumption that demand is a nonstationary Poisson process corresponds to choosing either the model setting of a constant VMR of 1 (this is the default and is currently used by the USAF) or the choice of the Vari-METRIC option with the shape parameter set to a very large value (low dispersion of the gamma prior distribution). The Vari-METRIC option handles the contribution of depot backorders to the base resupply pipeline more precisely than the constant VMR option and results in a better approximation of the PDF of base backorders.

Choosing a constant VMR larger than 1 or choosing the Vari-METRIC option with a small value for the shape parameter (large dispersion of the gamma prior distribution) results in modeling the demand in a time interval with a negative binomial distribution. Again, the Vari-METRIC option models the depot backorder portion of the base pipeline more precisely.

3.2.2 Backorder Probability Distributions

In Chapter 2, the first stage of estimating aircraft availability was to compute each item's expected backorders at the base, for which we needed the probability distribution of base backorders. Since we were assuming that the base ordered stock from the depot according to an $(s - 1, s)$ inventory policy, we had $BO = x - s$ for $x > s$, where BO was the number of base backorders, x was the number in the base resupply pipeline, s was the constant base inventory position, and both x and BO were random variables. From this equation, we saw that to compute the probability distribution for backorders, it sufficed to compute the distribution for the number in the base resupply pipeline x .

We then approximated the distribution of x in one of two ways. One was to approximate it with a Poisson distribution with mean equal to the sum of the means of the three components of the base resupply pipeline, which simplified the problem by allowing us to work with expected values alone, but introduced

some error because depot backorders do not usually have a Poisson distribution. The other was to use the Vari-METRIC technique, which uses a negative binomial for approximating the distribution. This introduces considerably less error in cases in which we have resupply from a depot with a finite, positive stock level [Slay].

Under dynamic conditions, we still focus on aircraft availability as our measure of supply performance, but — as noted earlier — we are now interested in availability on more than one day. We introduce one further complication: under dynamic conditions, we allow for a demand at the base to be satisfied by using a part from a grounded aircraft (cannibalization) — this is a policy likely to be used in war. Availability under cannibalization must be computed from the distribution of the number of aircraft NMCS for each item, as we will see in Chapter 4, rather than from each item's expected backorders.

However, both the distribution of the number of aircraft NMCS for an item and the item's expected backorders can be computed if we know the distribution for base backorders, so — as in the steady-state case — the problem of computing availability can be solved if we can compute that distribution (as a function of time). As in the steady-state case, we compute this distribution by first computing the probability distribution for the base resupply pipeline. The process is now made somewhat more complicated, since backorders and all components of the base resupply pipeline are nonstationary stochastic processes rather than ordinary random variables, so that the distributions we compute are time-dependent.

Sections 3.3 through 3.5 describe the dynamic pipeline methodology and how the model considers the effects of dynamic resupply times and FHPs. We explain how the model handles the transient effects, explicitly considering the lag between the actual component failures and the time they are felt in the supply system. We present the model algorithms here; for an example, see Appendix B.

Since the methodology is complex, we start by describing how to estimate pipeline distributions when resupply (repair, shipping, and procurement) times do not change over time, but where the demand process does. (The example in Appendix B shows flying hours as the driver for the demand process, but again, this is not essential.) Next we describe how to estimate pipeline distributions when those resupply times do change; we describe how the model handles the transition from steady-state to dynamic conditions. Finally, in Section 3.6, we present extensions to the calculation of pipeline distributions that address suspended resupply, SRU pipelines, days of warning, and multiple bases.

3.3 BACKORDERS WITH CONSTANT RESUPPLY TIMES

3.3.1 INTRODUCTION

To compute availability for time T — say day 30 — the model *looks back* to determine what previous conditions affect day 30. As discussed above, the key to availability is backorders, which are determined by the behavior of pipelines.

In the deterministic case, to compute how much is in a resupply pipeline today, we need to consider the earliest time that anything now in the pipeline could have entered. With constant resupply times, this “earliest time” is just a resupply time ago. For instance, to estimate what is in the order and ship pipeline on day 30, we look back an OST (say 10 days) from that day (to the end of day 20) and estimate what has failed in that period. A replacement for anything that has failed since day 20 could not have arrived from the depot and so is still in the pipeline. A replacement for anything that failed on day 20 or earlier has been received by the base by day 30 and so is out of the pipeline, unless it was backordered (by the depot) as of day 20.

In the general case, the quantity of items in the pipeline on day 30 is a random variable, *so we cannot compute how much is in the pipeline* on that day. Nonetheless, focusing on what happened in the period that started a resupply time ago is central to our treatment of dynamic pipelines.

3.3.2 Assumptions and Definitions

We assume that items fail according to a nonhomogeneous Poisson process with intensity $\lambda(t)$ on day t , where the process begins on day zero (there is no loss of generality in assuming this). We assume for now that the item is an LRU with no SRUs so that we can ignore the contribution of SRU backorders to the LRU base repair pipeline and so that there is a simple relationship between flying hours and mean failures. We will discuss SRU effects in Section 3.6.2.

For now, our inventory system consists of a single base and a single depot, and the base either repairs a failed item or ships it to the depot.² We further assume that the probability that an item that fails at time t is sent back to the depot is $NRTS(t)$, and that BRT and DRT are constant (DRT includes the retrograde time to ship the carcass from base to depot). We assume that the classification of an item either as base repairable or as NRTS is instantaneous and coincides with the time that the failure of the LRU results in a demand on base supply. We assume, for the moment, that the depot can satisfy all demands through repair; i.e., there are no condemnations.

Denote the base stock level by s , the depot stockage level by s_0 , and let the stockage policy at the base be $(s - 1, s)$. We assume that whenever a base orders

²We treat multiple bases later in Section 3.6.5 (also see Chapter 2); the assumption of a single base and depot may seem a bit artificial, but it simplifies the notation.

a part from the depot, the broken carcass is immediately shipped back to the depot to be repaired. Thus the depot will repair the returned carcass to replace the item issued to the base.

An item in the base resupply pipeline is either in base repair, in the order and ship pipeline, or backordered at the depot (the demand will eventually be filled by a serviceable item from depot repair). All of these three components of the base resupply pipeline are stochastic processes, and the quantity in any resupply pipeline at time T is a random variable.

Let OST be the constant (for now) order and ship time, and use the following notation for the resupply pipeline random variables:

$$\begin{aligned}
 BRpipeRV_T &= \text{number of items in base repair at time } T \\
 OSpiceRV_T &= \text{number of depot demands in the interval } (T - OST, T) \\
 DBORV_{T-OST} &= \text{number of depot backorders existing at time } T - OST \\
 DRpipeRV_T &= \text{number of items in depot repair pipeline at time } T \\
 BpipeRV_T &= \text{total number due in to the base at time } T \\
 BORV_T &= \text{number of base backorders at time } T.
 \end{aligned}$$

Then

$$\begin{aligned}
 BpipeRV_T &= BRpipeRV_T + OSpiceRV_T + DBORV_{T-OST}; \\
 BORV_T &= BpipeRV_T - s, \text{ if } BpipeRV_T > s; \text{ or} \\
 BORV_T &= 0, \text{ if } BpipeRV_T \leq s.
 \end{aligned}
 \tag{Eq. 3-1}$$

3.3.3 The Probability Distribution of Base Backorders

Our goal is to compute the probability distribution function of $BORV_T$, which we can compute from that of $BpipeRV_T$ (both functions of time.) Throughout this section, we continue to assume that the base repair, order and ship, and depot repair times are constant. We start by showing that the three random variables on the right-hand side of Equation 3-1 are mutually independent. Our development follows that of [King, Kruse, Simon].

From our assumption that demand was a nonhomogeneous Poisson process with intensity $\lambda(t)$, from the definition of $NRTS(t)$, and from the splitting theorem for nonhomogeneous Poisson processes [Kotkin], we conclude that the demands on base repair and the demands on the depot are independent nonhomogeneous Poisson processes with intensities $\lambda(t)[1 - NRTS(t)]$ and $\lambda(t) NRTS(t)$, respectively.

Let $X(t)$ be the number of demands on base repair in the time interval $[0, t]$ and $Y(u)$ be the number of demands on the depot in $[0, u]$. Since the demands on the base and demands on the depot are independent stochastic processes, the random vectors $[X(T), X(T - BRT)]$ and $[Y(T), Y(T - OST), Y(T - OST - DRT)]$ are independent. This follows directly from the definition of independent stochastic processes. Since functions of two independent random vectors are independent [Papoulis], $X(T) - X(T - BRT)$ is independent of $[Y(T) - Y(T - OST), Y(T - OST) - Y(T - OST - DRT) - s_0]$. We know that depot demand is a process with independent increments, so $Y(T) - Y(T - OST)$ is independent of $Y(T - OST) - Y(T - OST - DRT) - s_0$. Therefore the three random variables $X(T) - X(T - BRT)$, $Y(T) - Y(T - OST)$, and $Y(T - OST) - Y(T - OST - DRT) - s_0$ are mutually (not just pair-wise) independent.

Now $X(T) - X(T - BRT)$ is just the number of demands on base repair in the time interval $(T - BRT, T)$. Since any unit that entered base repair before $T - BRT$ has emerged by T , and any unit that has entered base repair after $T - BRT$ has not, $X(T) - X(T - BRT)$ is the number of units in the base repair pipeline at time T , $BRpipeRV_T$. Similarly, $Y(T) - Y(T - OST)$ is the number of units in the order and ship pipeline at T , $OSpipeRV_T$, and $Y(T - OST) - Y(T - OST - DRT)$ is the number of units in the depot repair pipeline at time $T - OST$, $DRpipeRV_{T-OST}$. The number of depot backorders at $T - OST$ is just $DBORV_{T-OST} = DRpipeRV_{T-OST} - s_0 = Y(T - OST) - Y(T - OST - DRT) - s_0$. Therefore $BRpipeRV_T$, $OSpipeRV_T$, and $DBORV_{T-OST}$ are mutually independent.

In obtaining the probability distributions of these random variables, we use the notation $\Pr(E)$ for the probability of an event E . As we observed, demands on base repair are a nonhomogeneous Poisson process with intensity $\lambda(t)[1 - NRTS(t)]$. Since $BRpipeRV_T$ is the number of demands on base repair in $[T - BRT, T]$, $BRpipeRV_T$ is a Poisson random variable with mean and variance equal to

$$\Phi(T) = \int_{T-BRT}^T \lambda(t)[1 - NRTS(t)]dt. \quad [\text{Eq. 3-2}]$$

(This follows directly from the definition of the mean-value function of a nonhomogeneous Poisson process presented earlier in Section 3.2.1.)

Since demands on the depot are a nonhomogeneous Poisson process with intensity $\lambda(t) NRTS(t)$, and since $OSpipeRV_T$ is just the depot demands in $[T - OST, T]$, we see that $OSpipeRV_T$ is a Poisson random variable with mean and variance equal to

$$\Psi(T) = \int_{T-OST}^T \lambda(t) NRTS(t)dt. \quad [\text{Eq. 3-3}]$$

To obtain the distribution of $DBORV_{T-OST}$, we need only compute the distribution of the depot repair pipeline at $T - OST$, for the probability that $DBORV_{T-OST} = k$ is just the probability that $DRpipeRV_{T-OST} = s_0 + k$ for $k > 0$ (recall that s_0 is the constant depot stock level) and the probability that $DBORV_{T-OST} = 0$ is $\sum_{j=0}^{s_0} \Pr(DRpipeRV_{T-OST} = j)$. But $DRpipeRV_{T-OST}$ is just the

number of depot demands in $(T - OST - DRT, T - OST)$, so by the fact that demands on the depot are a nonhomogeneous Poisson process with intensity $\lambda(t) NRTS(t)$, we see that $DRpipeRV_{T-OST}$ is a Poisson random variable with mean and variance

$$\Omega(T) = \int_{T-OST-DRT}^{T-OST} \lambda(t) NRTS(t) dt. \quad [\text{Eq. 3-4}]$$

Therefore the distribution of $DBORV_T$ is given by

$$\Pr(DBORV_T = 0) = \sum_{j=0}^{s_0} \frac{[\Omega(T)]^j e^{-\Omega(T)}}{j!}; \quad [\text{Eq. 3-5}]$$

$$\Pr(DBORV_T = k) = \frac{[\Omega(T)]^{s_0+k} e^{-\Omega(T)}}{(s_0 + k)!}, \text{ for } k > 0. \quad [\text{Eq. 3-6}]$$

We see that the distribution of $DBORV_{T-OST}$ is a spike at zero plus the tail of the distribution for the depot repair pipeline.

Since we can compute the distributions of all of the random variables on the right side in Equation 3-1, and since those random variables are independent, we could compute the distribution of the base resupply pipeline $BpipeRV_T$ by convolving the distributions for the three components of that pipeline. However, this convolution is computationally cumbersome.

Therefore we take a different approach. We recall from Chapter 2 that a negative binomial distribution is specified by two parameters, the mean μ and variance σ^2 , and we approximate the base resupply distribution with a negative binomial distribution having the same mean and variance. We have already found the mean and variance of $BRpipeRV_T$ and $OSpipeRV_T$; we can compute the mean and variance of the distribution for $DBORV_{T-OST}$ directly from its distribution:

$$E(DBORV_{T-OST}) = \sum_{k=1}^{\infty} k \frac{[\Omega(T)]^{s_0+k} e^{-\Omega(T)}}{(s_0 + k)!} \quad [\text{Eq. 3-7}]$$

$$E[(DBORV_{T-OST})^2] = \sum_{k=1}^{\infty} k^2 \frac{[\Omega(T)]^{s_0+k} e^{-\Omega(T)}}{(s_0 + k)!}; \quad [\text{Eq. 3-8}]$$

and

$$\text{Var}(DBORV_{T-OST}) = E[(DBORV_{T-OST})^2] - [E(DBORV_{T-OST})]^2. \quad [\text{Eq. 3-9}]$$

We also have shown that the three components of the base resupply pipeline at time T are independent random variables. Therefore we have

$$E(BpipeRV_T) = E(BRpipeRV_T) + E(OSpipeRV_T) + E(DBORV_{T-OST}) \quad [\text{Eq. 3-10}]$$

and

$$\text{Var}(B\text{pipe}RV_T) = \text{Var}(BR\text{pipe}RV_T) + \text{Var}(OS\text{pipe}RV_T) + \text{Var}(DBORV_{T-OST}). \quad [\text{Eq. 3-11}]$$

Note that although we have found the mean and variance of the distribution of the base resupply pipeline at time T , we do not know the form of that density.

In the model option of a constant VMR of 1, we approximate the distribution of the base resupply pipeline with a Poisson distribution with mean given by Equation 3-10. This introduces some error: unless the depot stock level is zero, the distribution of depot backorders will not be Poisson, so the distribution of the base resupply pipeline cannot be Poisson either. If we use the Vari-METRIC option (with gamma prior distribution nearly concentrated at a point, so that demand on the base is nonstationary Poisson), then we can get a better approximation of the base resupply pipeline distribution since this option explicitly handles the variance of depot backorders.

To approximate the base resupply pipeline distribution with a negative binomial density, let

$$\mu(T) = E(B\text{pipe}RV_T), \quad [\text{Eq. 3-12}]$$

$$\sigma^2(T) = \text{Var}(B\text{pipe}RV_T), \quad [\text{Eq. 3-13}]$$

and

$$\text{VMR}(T) = \frac{\sigma^2(T)}{\mu(T)}. \quad [\text{Eq. 3-14}]$$

With this notation, our approximation to the distribution of the base resupply pipeline at time T is given by

$$\Pr(B\text{pipe}RV_T = k) = \frac{\Gamma\left[k + \frac{\mu(T)}{\text{VMR}(T) - 1}\right]}{k! \Gamma\left[\frac{\mu(T)}{\text{VMR}(T) - 1}\right]} \left[\frac{\text{VMR}(T) - 1}{\text{VMR}(T)}\right]^k \text{VMR}(T)^{\frac{-\mu(T)}{\text{VMR}(T) - 1}}, \text{VMR}(T) \neq 1,$$

and

$$\Pr(B\text{pipe}RV_T = k) = \frac{e^{-\mu(T)} \mu(T)^k}{k!}, \text{ for } \text{VMR}(T) = 1, \quad [\text{Eq. 3-15}]$$

where Γ denotes the gamma function. Finally

$$\Pr(BORV_T = k) = \Pr(BpipeRV_T = s + k), \text{ for } k > 0, \quad [\text{Eq. 3-16}]$$

$$\Pr(BORV_T = 0) = \sum_{k=0}^s \Pr(BpipeRV_T = k) \quad [\text{Eq. 3-17}]$$

is the distribution function for base backorders that we were seeking.

Notice that none of the ideas developed so far in this section depend upon the demand intensities being driven by a weapon-system program, such as flying hours or sorties. For instance, we could use this method if we were estimating the intensity of future demand by using a trend-sensitive forecast based on historical demand. After a brief digression to cover computational methods, we will focus on the case in which demand is driven by a weapon-system program.

3.3.4 Computational Methods

We detour briefly here from the conceptual development to explain how the probability distribution for base backorders can be conveniently computed — not just as a function of time, but as a function of the stock levels as well — using a recursive process. We will frequently refer to the concepts developed previously, and in Appendix B we will use this section's results in a numerical example.

In Section 3.3.3, we held the stock levels fixed and focused on finding the time-dependent probability distribution of backorders for those levels. Now we will fix the time T and allow the stock levels s and s_0 to vary. To find item base and depot levels to maximize aircraft availability, the ASM evaluates the effect of various values of those levels for each item.

Consider a particular item. Of the three terms in Equation 3-10, only the last one depends on the depot stock level s_0 , and the same is true for Equation 3-11. Therefore to compute the change in the base resupply pipeline probability distribution due to adding a spare to the item's depot stock level, it suffices to compute the change in expected depot backorders and in the variance of depot backorders.

To emphasize the dependence of the mean and variance of depot backorders on s_0 , we will denote their values when the depot stock level is s_0 by $DEBO(s_0, T)$ and $VDBO(s_0, T)$, respectively. Let $P(k) = \Pr(DRpipeRV_T = k)$. When $s_0 = 0$, we have

$$\Pr(DBORV_T = k) = \Pr(DRpipeRV_T = k) = P(k) \text{ for } k \geq 0,$$

so that

$$DEBO(0, T) = E(DRpipeRV_T) \quad [\text{Eq. 3-18}]$$

and

$$VDBO(0, T) = \text{Var}(DRpipeRV_T). \quad [\text{Eq. 3-19}]$$

The right-hand sides of Equations 3-18 and 3-19 are known from the work laid out in Section 3.3.3 on the distribution function of base backorders.

We will develop recursion formulas for $DEBO(s_0, T)$ and $VDBO(s_0, T)$.

For any s_0 , we have

$$\begin{aligned} DEBO(s_0, T) &= \sum_{k=1}^{\infty} kP(s_0 + k) \\ &= P(s_0 + 1) + 2P(s_0 + 2) + 3P(s_0 + 3) + \dots \\ &= P(s_0 + 1) + P(s_0 + 2) + P(s_0 + 3) + \dots \\ &\quad P(s_0 + 2) + P(s_0 + 3) + \dots \\ &\quad P(s_0 + 3) + \dots \\ &\quad \dots \end{aligned} \quad [\text{Eq. 3-20}]$$

$$= \sum_{k=1}^{\infty} \Pr(DRpipeRV_T \geq s_0 + k). \quad [\text{Eq. 3-21}]$$

[We have gone from summing the triangular array above by columns (in Equation 3-20) to summing by rows (in Equation 3-21).]

Therefore we see that

$$\begin{aligned} &DEBO(s_0 + 1, T) - DEBO(s_0, T) \\ &= \sum_{k=1}^{\infty} \Pr(DRpipeRV_T \geq s_0 + k + 1) - \sum_{k=1}^{\infty} \Pr(DRpipeRV_T \geq s_0 + k) \\ &= \sum_{k=2}^{\infty} \Pr(DRpipeRV_T \geq s_0 + k) - \sum_{k=1}^{\infty} \Pr(DRpipeRV_T \geq s_0 + k) \\ &= -\Pr(DRpipeRV_T \geq s_0 + 1) \\ &= -\left[1 - \sum_{k=0}^{s_0} P(k)\right]. \end{aligned} \quad [\text{Eq. 3-22}]$$

That is, the reduction in expected backorders resulting from a change in stock level from s_0 to $s_0 + 1$ is $\left[1 - \sum_{k=0}^{s_0} P(k)\right]$. Note that as the stock level s_0 grows, the reduction in expected backorders resulting from an additional spare tends monotonically to zero (this is the “diminishing return” of adding spares). Because the decrease in expected backorders for each additional spare is a monotonically decreasing function of the number of spares, the number of expected backorders itself must be a decreasing, convex function of the number of spares.

By combining the recursion formula in Equation 3-22 with Equation 3-18, we can compute expected depot backorders for each stock level s_0 .

Next we show how to compute the variance of depot backorders. We will obtain the change in the variance of depot backorders resulting from an additional spare by computing the change in the second moment of depot backorders. Denote that second moment by $DE2BO(s_0, T)$. We have

$$DE2BO(s_0, T) = E\{[DBORV_T(s_0, T)]^2\} = \sum_{k=1}^{\infty} k^2 P(s_0 + k),$$

so that

$$\begin{aligned} DE2BO(s_0 + 1, T) - DE2BO(s_0, T) &= \sum_{k=1}^{\infty} k^2 P(s_0 + k + 1) - \sum_{k=1}^{\infty} k^2 P(s_0 + k) \\ &= \sum_{k=2}^{\infty} (k-1)^2 P(s_0 + k) - \sum_{k=1}^{\infty} k^2 P(s_0 + k) \\ &= \left[\sum_{k=2}^{\infty} (-2k+1) P(s_0 + k) \right] - P(s_0 + 1) \\ &= \left[-2 \sum_{k=2}^{\infty} k P(s_0 + k) + \sum_{k=2}^{\infty} P(s_0 + k) \right] - 2P(s_0 + 1) + P(s_0 + 1) \\ &= -2 \sum_{k=1}^{\infty} k P(s_0 + k) + \sum_{k=1}^{\infty} P(s_0 + k) \\ &= -2DEBO(s_0, T) + \left[1 - \sum_{k=0}^{s_0} P(k) \right]. \end{aligned}$$

But by Equation 3-22, the last term in brackets is just $DEBO(s_0, T) - DEBO(s_0 + 1, T)$,

so that

$$DE2BO(s_0 + 1, T) - DE2BO(s_0, T) = -2DEBO(s_0, T) + DEBO(s_0, T) - DEBO(s_0 + 1, T)$$

or

$$DE2BO(s_0 + 1, T) = DE2BO(s_0, T) - DEBO(s_0, T) - DEBO(s_0 + 1, T). \quad [\text{Eq. 3-23}]$$

From Equations 3-18 and 3-19, from the relation

$$DE2BO(0, T) = VDBO(0, T) + DEBO(0, T)^2 \quad [\text{Eq. 3-24}]$$

(which follows from the definition of variance) and the recursion formulas in Equations 3-22 and 3-23, we can compute the second moment for each depot stock level s_0 . We can then compute the variance for any particular depot stock level s_0 as

$$VDBO(s_0, T) = DE2BO(s_0, T) - DEBO(s_0, T)^2. \quad [\text{Eq. 3-25}]$$

Once we have computed the mean and variance of depot backorders for a particular value of s_0 , we use Equations 3-10 and 3-11 to estimate the mean and variance of the probability density for the base resupply pipeline. We then proceed to approximate the PDF of the base resupply pipeline as we did in Section 3.3.3, using the first part of Equation 3-15.

From the probability distribution of the base resupply pipeline and Equations 3-16 and 3-17, we can compute the distribution for base backorders as a

function of T , s_0 , and s . As noted at the beginning of this chapter, doing so allows us to compute our item's contribution to aircraft availability.

If the item is one for which cannibalization is not allowed, we need only compute expected base backorders for each combination of T , s_0 , and s , rather than the entire distribution function for base backorders, as we see from Equation 2-20. In this case there is a shortcut, based on the following observation: the technique developed for expected depot backorders can also be applied to compute the change in the expected base backorders due to adding an additional spare at the base (nothing in the mathematics depends on the fact that we were considering the *depot* resupply pipeline and *depot* backorders).

Since we know that expected base backorders with $s = 0$ is just the expected base resupply pipeline for any given depot stock level s_0 , we can use the technique we have just developed to compute the change in expected base backorders as we increment the base stock level s . Therefore we can determine expected base backorders as a function of T , s_0 , and s .

3.4 FLYING-HOUR-DRIVEN DEMANDS AND CONSTANT RESUPPLY TIMES

We now return to our conceptual development, starting with a brief discussion of how to apply the ideas of Section 3.3.3 to the case where demands are driven by flying hours. In Appendix B, we give a numerical example to illustrate the computation of time-dependent pipeline means and variances under a scenario typical of those to which the model is applied. As explained above, these means and variances are then used to compute the parameters of a time-dependent negative binomial distribution used to approximate the probability distribution of the base resupply pipeline.

Consider an LRU that has no subassemblies, and assume constant resupply times. Let $FF(t)$ and $FHP(t)$, respectively, be the item's mean failures per flying hour (of that type of aircraft) on day t . Let QPA and FAP be the quantity per aircraft and future application percentage, as defined in Chapter 2.

Assume that the item's demands are described by a nonhomogeneous Poisson process with intensity on day t given by

$$\lambda(t) = FF(t) \times FHP(t) \times QPA \times FAP \quad [\text{Eq. 3-26}]$$

and that the NRTS rate for the item, $NRTS(t)$, is also constant throughout any given day.

Denote the time-dependent means and variances of the pipelines at time T by

$$BRpipe(T) = E(BRpipeRV_T); \quad VBRpipe(T) = Var(BRpipeRV_T);$$

$$\begin{aligned}
OSpipe(T) &= E(OSpipeRV_T); & VOSpipe(T) &= Var(OSpipeRV_T); \\
DRpipe(T) &= E(DRpipeRV_T); & VDRpipe(T) &= Var(DRpipeRV_T); \\
DEBO(T) &= E(DBORV_T); & VDBO(T) &= Var(DRpipeRV_T); \\
Bpipe(T) &= E(BpipeRV_T); & VBpipe(T) &= Var(BpipeRV_T).
\end{aligned}$$

We begin by computing the mean of the base repair pipeline, $BRpipe(T)$, for an item on day T of the scenario. From Equation 3-2, we have

$$BRpipe(T) = \Phi(T) = \int_{T-BRT}^T \lambda(t)[1 - NRTS(t)]dt.$$

Since both the intensity of the demand process and the NRTS rate are step functions that are constant throughout any given day, the above integral reduces to the sum

$$BRpipe(T) = \sum_{k=T-BRT+1}^T \lambda(k)[1 - NRTS(k)]. \quad [\text{Eq. 3-27}]$$

Because the base repair pipeline was a nonhomogeneous Poisson process, we have $VBRpipe(T) = BRpipe(T)$.

By analogous reasoning, Equation 3-3 for the mean (and variance) of the $OSpipe$ on day T reduces to

$$OSpipe(T) = \Psi(T) = \sum_{k=T-OST+1}^T \lambda(k) NRTS(k). \quad [\text{Eq. 3-28}]$$

To obtain the mean and variance of $DBORV_{T-OST}$, we proceed as we did for the other two pipelines. Equation 3-4 for the mean and variance of the $DRpipe$ at time $T - OST$ reduces to

$$DRpipe(T - OST) = \Omega(T) = \sum_{k=T-OST-DRT+1}^{T-OST} \lambda(k) NRTS(k). \quad [\text{Eq. 3-29}]$$

From this, we compute $DEBO(T - OST)$ and $VDBO(T - OST)$, using Equations 3-7 and 3-9, respectively.

Recasting Equations 3-10 and 3-11 for the mean and variance of the base resupply pipeline in our current notation, we have

$$Bpipe(T) = BRpipe(T) + OSpipe(T) + DBO(T - OST) \quad [\text{Eq. 3-30}]$$

and

$$VBpipe(T) = VBRpipe(T) + VOSpipe(T) + VDBO(T - OST). \quad [\text{Eq. 3-31}]$$

As explained earlier, we then approximate the PDF of the base resupply pipeline at time T by a negative binomial distribution with mean $Bpipe(T)$ and variance $VBpipe(T)$, or by a Poisson distribution with mean $Bpipe(T)$, depending

on the VMR option selected in the model. (See Equations 3-12 through 3-15.) Appendix B contains a detailed numerical example illustrating these computations.

3.5 BACKORDERS UNDER VARYING DETERMINISTIC RESUPPLY TIMES

Now suppose we have an LRU with resupply times (base repair, order and ship, or depot repair) that change between peace and war. What does this mean? What happens to the items already in resupply at the start of the war? And how does the variability of those resupply times affect our computation of base backorder statistics?

As in the case of constant resupply times, the first step in computing backorder statistics is to focus on computing the means and variances of the components of the base resupply pipeline. The key to understanding the behavior of resupply pipelines is the retrospective point of view used earlier: What is in a pipeline today depends on the rate at which items entered the pipeline in the past and on the (now variable) resupply time.

Consider once more the example with deterministic pipelines that we used at the beginning of Section 3.3. In that example, the order and ship time was a constant 10 days. To determine the number of items that were in the pipeline on day $T = 30$, we needed to count all of the items that entered the pipeline between day $T - OST = 20$ and day 30, since only items entering the pipeline during this period could still be in the pipeline.

If we now suppose that the OST is variable, we can no longer compute what is in the pipeline on day T simply by looking back an order and ship time, since there is no single order and ship time that will do. We will show that there is an earliest time that an item could have entered the order and ship pipeline without having emerged by time T . The difference between this **induction time**, as we will call it, and the **emergence time** will play the role that the order and ship time played in the case of constant resupply times.

For the rest of Section 3.5, resupply pipelines are again stochastic processes (by this, we mean that the number of units in any resupply pipeline at a time t is a random variable whose PDF depends on t ; we do not mean that the resupply time is random). We continue to assume that the demand process is a nonhomogeneous Poisson process.

3.5.1 Resupply Processes with Deterministic Variable Resupply Times

We will assume for now that all resupply processes have deterministic resupply times and are "first in, first out" (FIFO). By "deterministic resupply

times," we mean that all simultaneous demands for a particular component with identical sources of resupply have the same resupply time. In other words, all demands on base supply for a particular component occurring on the same day will have the same known resupply time. A resupply process is FIFO if the fact that item one emerges from resupply before item two implies that item one entered the resupply process before item two did.

Note that if a resupply process is FIFO, the number of units in resupply at a time t is still a random variable; it is only the resupply *times* that are deterministic. Because resupply times can vary from one day to another, "deterministic resupply time" does not necessarily mean "constant."

Since we are assuming that our resupply times are deterministic, the time at which an item emerges from resupply uniquely defines the time when it entered. Therefore for a resupply process with a deterministic resupply time, we can define the function

$I(t)$ = induction time (time at which the resupply action is initiated)
for an item that emerges from resupply at time t .

We see that a deterministic process is FIFO if and only if

$$I(t_1) < I(t_2) \text{ for all times } t_1, t_2 \text{ with } t_1 < t_2.$$

The induction time function splits history into two segments. At any time t , all items inducted before $I(t)$ will have emerged from resupply; all items inducted after $I(t)$ will not have.

The resupply time is the difference between t and $I(t)$. Since the model looks back in time to estimate availability, it computes resupply time on the basis of when something is inducted — when it enters the resupply process. *An item that emerges at time t has resupply time*

$$R(t) = t - I(t). \quad [\text{Eq. 3-32}]$$

We will see how the model computes $I(t)$ later, when we discuss the way the model handles the behavior of resupply processes under the peace-to-war transition.

3.5.2 The Probability Distribution of Base Backorders with Varying Resupply Times

As in the case of constant resupply times, we must compute the probability distribution of base backorders. Thus we must first compute the means and variances of the segments of the base resupply pipeline. Our development is almost identical to that in the constant resupply time case; the only difference is that constant resupply times in the equations are replaced by resupply times that

are functions of the emergence time. We will present only those parts of the earlier development that have changed.

We denote the time-dependent resupply times by

$OST(T)$ = order and ship time for an item emerging at time T ,³

$BRT(T)$ = base repair time for an item emerging at time T ;

$DRT(T)$ = depot repair time for an item emerging at time T .

With this notation, we have the following analog of Equation 3-1:

$$BpipeRV_T = OSpiceRV_T + BRpipeRV_T + DEBO_{T-OST(T)}. \quad [\text{Eq. 3-33}]$$

That is, the base resupply pipeline at time T (a random variable) is the sum of the order and ship pipeline at time T , the base repair pipeline at time T , and the depot backorders an order and ship time previous to T . We denote that time by $OST(T)$ to show its dependence on the emergence time. We see that $T - OST(T)$ is playing the role of $I(t)$ in Equation 3-32.

As we saw earlier with Equation 3-1, the three terms on the right-hand side of Equation 3-33 are independent random variables. The only changes required in our earlier argument are that we view the order and ship pipeline at time T as being the number of demands on the depot in the time interval $[T - OST(T), T]$, the base repair pipeline as the number of demands on the base in the time interval $[T - BRT(T), T]$, and the depot repair pipeline as the number of demands on the depot in the time interval $[T - OST(T) - DRT(T - OST), T - OST(T)]$. Because we are still assuming that demand on the base is a nonhomogeneous Poisson process, the argument for the independence of the three random variables is unchanged. The nonhomogeneous Poisson assumption also tells us that we have the analogs of Equations 3-2, 3-3, and 3-4:

$$\Phi(T) = \int_{T-BRT(T)}^T \lambda(t) [1 - NRTS(t)] dt; \quad [\text{Eq. 3-34}]$$

$$\Psi(T) = \int_{T-OST(T)}^T \lambda(t) NRTS(t) dt; \quad [\text{Eq. 3-35}]$$

$$\Omega(T) = \int_{T-OST(T)-DRT(T-OST(T))}^{T-OST(T)} \lambda(t) NRTS(t) dt; \quad [\text{Eq. 3-36}]$$

and that $\Phi(T)$, $\Psi(T)$, and $\Omega(T)$, respectively, are both the means and variances of the base repair pipeline at time T , the order and ship pipeline at time T , and the depot repair pipeline at time T .

Knowing that the three components of the base resupply pipeline at time T are independent allows us to conclude that we can simply add the means and

³Here $OST(T)$ is a pure order and ship time containing no backorder delay.

variances of those components to obtain the mean and variance of the base resupply pipeline at time T , as in Equations 3-10 and 3-11.

The rest of our process for computing the distributions of the base resupply pipeline and of base backorders is unchanged from the case of constant resupply times (Equations 3-12 through 3-17), so we will not reproduce it here.

3.5.3 Flying-Hour-Driven Demands and Variable Resupply Times

As we did in Section 3.4, we now focus on the case where demands are driven by flying hours (Equation 3-26), and where flying hours and the NRTS rate are step functions that change only at the beginning of each day. We now assume that resupply times are also step functions that change only at the beginning of the day.

With these assumptions, Equations 3-34 through 3-36 for the means and variances of the base resupply pipeline components reduce from integrals to sums, as they did in Equations 3-27 through 3-29:

$$BRpipe(T) = \sum_{k=T-BRT(T)+1}^T \lambda(k)[1 - NRTS(k)]; \quad [\text{Eq. 3-37}]$$

$$OSpipe(T) = \Psi(T) = \sum_{k=T-OST(T)+1}^T \lambda(k) NRTS(k); \quad [\text{Eq. 3-38}]$$

$$DRpipe[T - OST(T)] = \Omega(T) = \sum_{k=T-OST(T)-DRT[T-OST(T)]+1}^{T-OST(T)} \lambda(k) NRTS(k). \quad [\text{Eq. 3-39}]$$

To use these equations, we must compute $BRT(T)$, $OST(T)$, and $DRT[T - OST(T)]$, which further specify resupply behavior. We assume that in peace, the resupply times have one constant value, and that after the war has been going on for some time, they have another constant value. During the first few days of the war, which we will call the transition period, we assume that resupply times are sufficiently well approximated by linear interpolation between their peacetime and wartime values so that the FIFO property of base repair is maintained.

For example, consider the base repair pipeline. Suppose that the war begins at $T = 0$, and that we have

$$BRT(T) = BRT_p, \text{ for } T < 0;$$

$$BRT(T) = BRT_w, \text{ for } T > BRT_w.$$

Then we assume that

$$BRT(T) = BRT_p - T \times \frac{(BRT_p - BRT_w)}{BRT_w}, \text{ for } 0 \leq T \leq BRT_w, \quad [\text{Eq. 3-40}]$$

where $BRT(T)$ is rounded to the nearest integer.

To illustrate, suppose that the peacetime BRT is 7 days and the wartime BRT is 3 days. Table 3-1 displays the model's BRT by day of the scenario.

Table 3-1.
Peace-to-War Transitional Base Repair Times by Scenario Day

Day t	-3	-2	-1	0	1	2	3	4	5	6
$BRT(t)$	7	7	7	7	6	4	3	3	3	3
$I(t)$	-10	-9	-8	-7	-5	-2	0	1	2	3

Using Equations 3-37 through 3-40, we can compute the means and variances of the components of the base resupply pipeline during the early days of the war, when the supply system is in transition. (For other days, we revert to the constant-resupply-time case treated earlier in Section 3.4.) Computing the probability distribution of base backorders proceeds as in the case of constant resupply times.

3.6 OTHER EXTENSIONS

For clarity's sake, we have so far simplified some of the hypotheses concerning the wartime scenario. Now we describe briefly some extensions of our theory. They represent straightforward modifications of the standard ASM treatment and of the dynamic theory discussed in this report and earlier [King].

3.6.1 Suspended Resupply

In many wartime or contingency situations, repair facilities may not be immediately operational after a deployment, or transportation vehicles may be dedicated to troop movement and unavailable for spares transport. In such cases, repair and resupply are suspended for some length of time. The examples so far have assumed that all types of resupply began on day zero (that is, no suspension of repair or transportation). While we will examine suspension of base repair only, the formulas for depot repair suspension and for order and ship suspension are analogous.

Suppose base repair is suspended at the start of the war though time ST , i.e., repair begins again on day $ST + 1$. This suspension has two effects, as shown in Table 3-2. The first is that no item in repair at the start of the war will finish repair until after the suspension is over; that is, the time an item spends in the repair pipeline increases by one day for each day of the repair suspension. During the suspension, $BRT(t)$ is best understood as the time an item has been in the

Table 3-2.
Base Repair Time with Suspension

Time	Base repair time
$t \leq 0$	$BRT(t) = BRT_p$
$0 < t \leq ST$	$BRT(t) = BRT_p + t$
$ST < t \leq BRT_w + ST$	$BRT(t) = (BRT_p + ST) - (t - ST) \times [(BRT_p - BRT_w)/BRT_w]$
$BRT_w + ST < t$	$BRT(t) = BRT_w$

pipeline,⁴ rather than a repair time. The second effect is that items broken in war do not go into repair until the suspension is over (i.e., all items broken during the suspension go into repair after suspension and come out of repair a BRT_w later. In between those two times (ST and $BRT_w + ST$), the time required to repair an item decreases in the same manner as in the no-suspension peace-to-war transition case described earlier [Equation 3-40]. Table 3-3 demonstrates that point by comparing a 3-day suspension starting on day 1 with our previous example (see Table 3-1) with no suspension. The BRT for peace and war equals 7 and 3 days, respectively, for both examples. Notice that the suspended case repair times decrease by the same rate (days 4, 5, and 6) as the case with no suspension during days 1, 2, and 3.

Table 3-3.
Peace-to-War Transitional Base Repair Times with Suspension

Day t	-1	0	1	2	3	4	5	6	7	8
$BRT(t)$ with suspension	7	7	8	9	10	9	7	6	3	3
$BRT(t)$ no suspension	7	7	6	4	3	3	3	3	3	3

3.6.2 Levels of Indenture

3.6.2.1 INTRODUCTION

By "indenture" we mean the relationship of a subassembly (lower indenture) to its parent assembly (LRU). With this terminology, an LRU is referred to as a first-indenture item, since it is installed directly on the aircraft. An SRU is a second-indenture or lower-indenture item. Thus "levels of indenture" refers to a hierarchical relationship between major assemblies and their subassemblies.

To this point, we have considered LRUs only. When an LRU has SRUs, base repair typically consists of identifying, removing, and replacing a failed SRU. The repair of the LRU may be delayed as a result of the time needed to obtain an

⁴Because we are only allowing t to take on integral values in the model, this time in the pipeline is approximate for some values of t .

SRU either from repair or from resupply. The model must estimate the time the LRU spends waiting for resupply of the SRU (which we term **SRU delay**) and the resulting degradation of aircraft availability. This delay can range from zero, if the SRU has plenty of spares, to the full resupply time for the SRU, if no spares are available. We consider only the effect of the SRU delay on base repair of the parent LRU; we assume no SRU delay on depot repair of the parent LRU.⁵

In Section 3.6.2, we show how the calculation of the mean and variance of the base resupply pipeline for an LRU at a time T changes to include the effect of SRU backorders (we must defer part of the explanation until Chapter 4, when we discuss cannibalization, since the model assumes that all SRUs are cannibalized if feasible). The theory that we have applied so far to derive the distribution of LRU backorders can also be used to derive the distribution of SRU backorders. However, since SRU failures are not detected until the LRU has undergone fault isolation, the failure process for the SRUs effectively lags the flying-hour program by the LRU repair time. The remaining steps in the calculation of the base backorder PDF are unchanged from earlier sections. We continue to assume that demand at the base is a nonhomogeneous Poisson process for LRUs, and now for SRUs as well. At the end of Section 3.6.2, we present a numerical example that illustrates the relationship between the time of an SRU failure and its effect on LRU repair.

3.6.2.2 APPROACH

Consider an LRU with a single SRU and assume constant base repair times. (Constant resupply times are not necessary here, but they simplify the notation.) Assume further that there is no stock for the SRU, so that we can see the full effect of the SRU delay time. Let BRT and $BRTLRLU$, respectively, represent the BRTs for a specific SRU and for its LRU parent. $BRTLRLU$ can be written as

$$BRTLRLU = BRTLRLU_{FIT} + BRTLRLU_{RAT}, \quad [\text{Eq. 3-41}]$$

where

$BRTLRLU_{FIT}$ = fault isolation time for the LRU, and
 $BRTLRLU_{RAT}$ = reassembly time for the LRU.

⁵The flow of depot repair is much more complex than that of base repair. In many cases, the standard DRT already includes ample allowance for any SRU delay. Frequently, in fact, a failed SRU is "job-routed," i.e., repaired as part of the repair of its parent LRU. Job-routing precludes any demand on supply for the SRU, unless the SRU is condemned.

The sequence of events in the LRU repair process is depicted in Figure 3-2. We assume that the reassembly time is negligible in comparison with the fault isolation time.⁶ Then $BRTLRLU = BRTLRLU_{FIT}$, approximately. Note from Figure 3-2 that SRU repair induction lags behind the corresponding LRU failure by the parent BRT. That is, the SRU failures occurred earlier, driven by flying hours at that time. However, those SRU failures are not discovered until after the LRU goes through a BRT. The SRU has an impact upon availability only at the point when the fault in the parent LRU has been isolated, and only if there are no SRU spares to repair the LRU.




Figure 3-2.
LRU/SRU Timeline

3.6.2.3 THE EFFECT OF SRU BACKORDERS

We now drop the assumptions of constant resupply times and no SRU stock. Denote the BRT for an SRU that emerges from repair at time T by $BRT(T)$, and the BRT for the parent LRU that emerges from repair at time T by $BRTNHA(T)$. Here, NHA means "next higher assembly." Assume that the SRU failure process is a nonhomogeneous Poisson process with intensity $\alpha(t)$ and postulate NRTS rate $NRTS(t)$. As with the LRU demand process considered earlier, the splitting theorem for nonhomogeneous Poisson processes tells us that the failure process for the SRU splits into two nonhomogeneous Poisson processes with intensities $[1 - NRTS(t)] \alpha(t)$ and $NRTS(t) \alpha(t)$. The first process is that of SRU demands on base repair, while the second is that of SRU demands on the depot.

⁶In fact, the reassembly time need only be nondynamic for the following results to apply. SRU failures are delayed by the LRU fault isolation time. However, the impact of SRU backorders on LRU backorders is also delayed by the LRU reassembly time. (This is exactly analogous to the way the impact of depot backorders on the base is delayed by the order and ship time.) Thus the total delay of the effect of SRU backorders is $BRTLRLU$. But it is the dynamics of the LRU fault isolation time that has a critical effect on LRU backorders. We assume that all of the dynamics of $BRTLRLU$ are in the fault isolation time and that the reassembly time is constant.

The SRU base repair pipeline at time T (a random variable) is the sum of the demands on base repair over the interval

$$[T - BRT(T), T].$$

Now an SRU demand on base repair at time T arose from a failure that occurred at a time $BRTNHA(T)$ ago, since $BRTNHA(T)$ is the fault isolation time required to discover that the SRU has failed. Similarly, an SRU demand on base repair at time $T - BRT(T)$ arose from a failure at a time $BRTNHA[T - BRT(T)]$ ago. Therefore the SRU demands on base repair during the time interval $[T - BRT(T), T]$ arose from SRU failures over the interval

$$\{T - BRT(T) - BRTNHA[T - BRT(T)], T - BRTNHA(T)\}. \quad [\text{Eq. 3-42}]$$

For brevity, let $T_1 = BRTNHA[T - BRT(T)]$ and $T_2 = BRTNHA(T)$.

From our assumption that the SRU failure process was nonhomogeneous Poisson with intensity $\alpha(t)$ and the above observation about the time-lagged relationship between SRU failures and SRU demands on base repair, we see that the SRU base repair pipeline at time T is a Poisson random variable with mean and variance

$$\Phi(T) = \int_{T-BRT(T)-T_1}^{T-T_2} [1 - NRTS(t)] \alpha(t) dt. \quad [\text{Eq. 3-43}]$$

The reader may wonder why this is not exactly analogous to the depot pipeline, with SRU failures over the interval

$$\{T - BRTNHA(T) - BRT[T - BRTNHA(T)], T - BRTNHA(T)\}.$$

That is, why do we not use the SRU's BRT a $BRTNHA$ ago? The answer is that we are assuming that most of the LRU's BRT is fault isolation time. If the situation were reversed and most of the LRU's BRT were reassembly time, then these pipelines would be analogous to the depot pipelines, where depot repair is followed by order and ship. In that case, the second form of the time interval would apply. However, modeling the main portion of $BRTNHA$ as preceding the SRU replacement is preferable for two reasons: first, when repair is expedited, this primarily shortens the scheduling delays, effectively shortening fault isolation time; second, fault isolation and removal of the SRU normally takes longer than reassembly.

Turning next to the order and ship pipeline for the SRU, we see that the number of items in that pipeline at time T is the number of SRU demands on the depot over the time interval $[T - OST(T), T]$. Reasoning as we did above, an SRU demand on the depot at time T arose from an SRU failure at time $T - BRTNHA(T)$, while a demand on the depot at time $T - OST(T)$ arose from

SRU failures at time $T - OST(T) - BRTNHA[T - OST(T)]$. Therefore the SRUs in the order and ship pipeline at time T arose from SRU failures over the time interval from

$$T - OST(T) - BRTNHA[T - OST(T)]$$

to

$$T - BRTNHA(T).$$

Letting

$$T_3 = BRTNHA[T - OST(T)],$$

we see that the SRU $OSpipe$ at time T is a Poisson random variable with mean and variance given by

$$\Psi(T) = \int_{T-OST(T)-T_3}^{T-T_2} NRTS(t) \alpha(t) dt. \quad [\text{Eq. 3-44}]$$

Now consider the SRU depot repair pipeline at time $T - OST(T)$. This determines SRU depot backorders at time T . The number of items in the depot repair pipeline at time T is the number of SRU depot demands over the time interval from $T - OST(T) - DRT[T - OST(T)] - BRTNHA\{T - OST(T) - DRT[T - OST(T)]\}$

to

$$T - OST(T) - BRTNHA[T - OST(T)].$$

Let

$$T_4 = BRTNHA\{T - OST(T) - DRT[T - OST(T)]\},$$

and

$$T_5 = BRTNHA[T - OST(T)];$$

then the SRU depot repair pipeline at time $T - OST(T)$ is a Poisson random variable with mean and variance

$$\Omega(T) = \int_{T-OST(T)-DRT[T-OST(T)]-T_4}^{T-OST(T)-T_5} NRTS(t) \alpha(t) dt. \quad [\text{Eq. 3-45}]$$

We can compute the PDF for SRU depot backorders at time $T - OST(T)$ from Equations 3-45, 3-5, and 3-6. Next, Equations 3-10 and 3-11 produce the mean and variance of the SRU base resupply pipeline at time T . As we did before, we use the Vari-METRIC approximation to compute the distributions of the SRU base resupply pipeline and the SRU base backorders at time T .

From the distribution of SRU base backorders, we can estimate the mean and variance of the **awaiting parts** (AWP) pipeline for the LRU. We defer an explanation of how this is done until Chapter 4, since the model uses the assumption that all SRUs are cannibalized in estimating the AWP pipeline mean and variance, and cannibalization is explained in Chapter 4.

The means and variances of the LRU's AWP pipeline are then added to the mean and variance, respectively, of the base resupply pipeline for the parent LRU at time T ; this is an approximation of the SRU effect because SRU backorders — and the part of the base resupply pipeline for the parent LRU that does not include SRU effects — are not independent processes. Finally, the model uses the Vari-METRIC method to estimate the distributions of the new base resupply pipeline (including SRU effects) and base backorders.

As already noted in Sections 3.4 and 3.5, the model assumes that the demand process, the NRTS rate, and the resupply times are day-by-day step functions: we omit the derivation of the above results for this special case because it is so similar to what has been presented before.

3.6.2.4 EXAMPLE

Since the derivation just presented is fairly complex, we now give an example illustrating how SRU failures have a time-lagged effect on the parent LRU. For simplicity, we assume constant resupply times.

Figure 3-3 shows the distinction between the time when the SRU fails and the time when it is inducted into maintenance. The figure assumes an impact point T on day 30 and constant LRU and SRU repair times of 10 and 6 days, respectively. Note that any SRU failure occurring on or before day 14 ($30 - 10 - 6$) has had time to go through the LRU and SRU repair process and so does not have an impact on the LRU. The first failure that can have an impact on the LRU is one occurring on day 15 (the top horizontal line in Figure 3-3 and the left side of Equation 3-42). (If no SRU spares were available, the SRU would still be undergoing repair and delaying the LRU.) The latest SRU failure that can have an impact on the LRU on day 30 would be one occurring on day 20. After that point, any SRU failures are not inducted into maintenance until after day 30 and so cannot have an impact on availability on day 30. That last point of impact is illustrated by the bottom horizontal line of Figure 3-3; it is the right side of Equation 3-42. Thus to determine the mean SRU pipeline, the model estimates mean SRU failures between the two extremes (day 15 and day 20), a period whose length equals the SRU repair time of 6 days (see the horizontal bar at the bottom of Figure 3-3).




Figure 3-3.
SRU Pipeline Computation

3.6.3 Condemnations

Consider an LRU with no SRUs, an LRU whose demands on the base form a nonhomogeneous Poisson process with intensity $\lambda(t)$. Thus far we have been assuming that any item that failed can be repaired, either by the base or by the depot; we now drop that assumption. We now interpret $NRTS(t)$ as the probability that a demand on the base at time t cannot be repaired by the base, including the possibility that it is condemned (either by the base or by the depot.) Let $ConPCT(t)$ be the probability of condemning a failed item at time t .

There are now three mutually exclusive events that can occur when there is a demand on the base at time t : either the item goes to base repair, it goes to depot repair, or it is condemned. The probability that the failure leads to a demand on base repair is still $1 - NRTS(t)$, as it was in the no-condemnations case. The probability that it leads to a demand on depot repair is now $NRTS(t) - ConPCT(t)$, and the probability that it results in a condemnation is $ConPCT(t)$.

An application of the splitting theorem for nonhomogeneous Poisson processes [Kotkin] tells us that these three probabilities split the base demand process into three independent nonhomogeneous Poisson processes with intensities $[1 - NRTS(t)]\lambda(t)$, $\lambda(t)[NRTS(t) - ConPCT(t)]$, and $\lambda(t)ConPCT(t)$. The first process is the demands on base repair, the second process is the demands on depot repair, and the third process is condemnations (which can be viewed as demands for procurement). Since demands on depot repair and condemnations

are independent nonhomogeneous Poisson processes with intensities $\lambda(t)[NRTS(t) - ConPCT(t)]$ and $\lambda(t)ConPCT(t)$, respectively, total demands on the depot (the sum of these two processes) is still a nonhomogeneous Poisson process with intensity $\lambda(t)[NRTS(t)]$, as it was in the no-condemnations case.

Because the three processes above are mutually independent, the total depot demand process, which is the sum of the latter two, is independent of the demands on base repair.

Since the demand process for base repair is unchanged (from the no-condemnations case), and the base repair pipeline at time T is still demands on base repair over the time interval $[T - BRT(T), T]$, the mean and variance of the base repair pipeline are still given by Equation 3-34.

The order and ship pipeline at time T is the total demands on the depot over the time interval $[T - OST(T), T]$, and the demand process on the depot is still the same as it was in the no-condemnations case, so the mean and variance of the pipeline are still given by Equation 3-35.

The depot repair pipeline at time $T - OST$ is the demands on depot repair over the time interval $[T - OST(T) - DRT[T - OST(T)], T - OST(T)]$. Because the depot repair demand process is nonhomogeneous Poisson with intensity $\lambda(t)[NRTS(t) - ConPCT(t)]$, we see that the mean and variance of the depot repair pipeline at time $T - OST(T)$ are both

$$\Omega(T) = \int_{T-OST(T)-DRT[T-OST(T)]}^{T-OST(T)} \lambda(t)[NRTS(t) - ConPCT(t)]dt. \quad [Eq. 3-46]$$

The condemnations at the depot at time $T - OST(T)$ are treated as another pipeline at the depot.⁷ Let PLT be the (constant) procurement lead-time for the LRU in days, with PLT representing the sum of production, administrative, and other processing times. The "condemnation pipeline" at time $T - OST(T)$ is the condemnations over the time interval $[T - OST(T) - PLT, T - OST(T)]$. Because condemnations are a nonhomogeneous Poisson process with intensity $NRTS(t)\lambda(t)ConPCT(t)$, the mean and variance of the condemnation pipeline at time $T - OST(T)$ are both

$$\Theta(T) = \int_{T-OST(T)-PLT}^{T-OST(T)} \lambda(t)NRTS(t)ConPCT(t)dt. \quad [Eq. 3-47]$$

Since the demands on depot repair and condemnations are independent stochastic processes, the depot repair pipeline and the condemnation pipeline at time $T - OST(T)$ are independent random variables. Therefore the total depot resupply pipeline at time $T - OST(T)$ is a Poisson random variable with mean and variance given by $\Omega(T) + \Theta(T)$. The PDF of depot backorders is computed from the PDF of the total depot resupply pipeline in the same way that we derived Equations 3-5 and 3-6, with $\Omega(T) + \Theta(T)$ playing the role that $\Omega(T)$ played in those earlier equations.

⁷ As in the steady-state case of Chapter 2, this is an idealization. Typically, procurement decisions are made periodically and not with an $(s - 1, s)$ discipline.

As in the case of constant resupply times and no condemnations, the base repair pipeline at time T , the order and ship pipeline at time T , and the depot backorders at time $T - OST(T)$ are mutually independent random variables. From this point on, the derivation of the PDF for *base* backorders remains unchanged from the derivation in the no-condemnation case.

If the LRU has SRUs, the probability distribution for SRU depot backorders is derived just as it was above for LRUs, and the effect of the SRU backorders on the LRU is handled just as it was in the case of no SRU condemnations.

When demands are driven by flying hours or other program units — and flying hours, NRTS rates, and condemnations are day-by-day step functions — Equations 3-46 and 3-47 reduce to sums. We omit these, since the transition from integrals to sums in this case is analogous to the transitions explained earlier in this chapter.

3.6.4 Days of Warning

Let N_w equal the number of days of warning before the start of the surge conflict.⁸ We model this situation by shifting the time dependence of the component characteristics by N_w days. For example, if the wartime base repair rate is 5 days and N_w is 3 days, the model interprets the 5-day base repair rate as beginning on day -2 (i.e., the last 3 days of peacetime base repairs are performed at the wartime rate). The other resupply times, as well as the NRTS and condemnation rates, are treated similarly. However, the flying hours and failure rates are not shifted.

3.6.5 Multiple Bases

Thus far, to simplify the mathematics, we have assumed that the supply system consists of one base and one depot. We now consider the case where several bases are resupplied from a single depot. The ASM assumes that all bases are uniform with respect to demand rates, resupply times, NRTS rates, and condemnation rates. Suppose there are N uniform bases, $N > 1$.

For base j at time T , denote the base repair pipeline by $BRpipeRV_j(T)$, the base order and ship pipeline by $OSpipeRV_j(T)$, depot backorders DI (owed) to the base by $DBORV_j(T)$, the total base resupply pipeline by $BpipeRV_j(T)$, and the base backorders by $BORV_j(T)$. Our goal is to compute the PDF of $BORV_j(T)$.

We know that base j 's resupply pipeline at time T is given by

$$BpipeRV_j(T) = BRpipeRV_j(T) + OSpipeRV_j(T) + DBORV_j[T - OST(T)]. \quad [\text{Eq. 3-48}]$$

⁸Here w means warning as opposed to war.

As before, the three terms on the right are independent random variables. Thus if we know their means and variances, we can sum the means and sum the variances to obtain the mean and variance, respectively, of the base resupply pipeline.

There is no change in the way we compute the mean and variance of the base repair and order and ship pipelines from the way we computed them in the single-base case; each base has its own base repair and order and ship processes (with identical parameters). However, the depot backorder segment requires a new treatment because it is not clear how to allocate depot backorders among bases.

We *do* know how to compute the distribution of *total* depot backorders — we get this from the distribution of the total depot resupply pipeline, as we did in our Section 3.6.3 on condemnations. Although the depot repair and condemnation pipelines now involve demands originating at all bases, the base of origin is immaterial to the calculation of the total depot resupply pipeline statistics.

We now show how to compute the mean and variance of depot backorders owed to base j ; we will use the following standard notation. When X and Y are random variables, we will use $P(X|Y)$, $E(X|Y)$, and $VAR(X|Y)$, respectively, to denote the (conditional) probability distribution of X given Y , the expectation of X given Y , and the variance of X given Y . We will use the following two results from probability theory [Ross]:

$$E(X) = E[E(X|Y)], \quad [\text{Eq. 3-49}]$$

and

$$VAR(X) = VAR[E(X|Y)] + E[VAR(X|Y)]. \quad [\text{Eq. 3-50}]$$

It can be shown under the uniform base assumption that the probability that any depot backorder represents an item due in to a particular base is $\frac{1}{N}$, and that the probabilities that any two successive depot backorders are owed to base j are independent. Thus we can view the occurrence of depot backorders as a sequence of Bernoulli trials [Feller, Vol. 1], where the probability of success, interpreted as the probability that a depot backorder belongs to base j , is $\frac{1}{N}$. This is an immediate consequence of a more general result in [Kotkin]; an earlier result for stationary demand processes is in [Kruse].

Let $DBORV(T)$ be total depot backorders at time T ; we know its mean and variance. With this notation, we have

$$\Pr[DBORV_j(T) = k | DBORV(T) = n] = \binom{n}{k} \left(\frac{1}{N}\right)^k \left(1 - \frac{1}{N}\right)^{n-k}.$$

In other words, the conditional distribution of depot backorders due in to base j given n total depot backorders is binomial with parameter $\frac{1}{N}$. Since the mean of a binomial distribution with parameters n and p is np [Feller, Vol. 1], we have

$$E[DBORV_j(T)|DBORV(T) = n] = \frac{n}{N} = \frac{DBORV(T)}{N},$$

and by Equation 3-49, we obtain

$$\begin{aligned} E[DBORV_j(T)] &= E\{E[DBORV_j(T)|DBORV(T)]\} \\ &= E\left[\frac{DBORV(T)}{N}\right] \\ &= \frac{1}{N}E[DBORV(T)]. \end{aligned} \quad [\text{Eq. 3-51}]$$

We have found the mean of the base resupply time for an arbitrary base; we now find the variance. Since the variance of a binomial distribution with parameters n and p is $np(1-p)$ [Feller, Vol. 1], we see that

$$\begin{aligned} \text{VAR}[DBORV_j(T)|DBORV(T) = n] &= n \times \frac{1}{N} \left(1 - \frac{1}{N}\right) \\ &= n \frac{N-1}{N^2} \\ &= \frac{N-1}{N^2} DBORV(T). \end{aligned} \quad [\text{Eq. 3-52}]$$

Using Equations 3-50, 3-51, and 3-52, we obtain

$$\begin{aligned} &\text{VAR}[DBORV_j(T)] \\ &= \text{VAR}\{E[DBORV_j(T)|DBORV(T)]\} + E\{\text{VAR}[DBORV_j(T)|DBORV(T)]\} \\ &= \text{VAR}\left[\frac{1}{N} DBORV(T)\right] + E\left[\frac{N-1}{N^2} DBORV(T)\right] \\ &= \frac{1}{N^2} \text{VAR}[DBORV(T)] + \frac{N-1}{N^2} E[DBORV(T)]. \end{aligned} \quad [\text{Eq. 3-53}]$$

Since we have found the mean and variance of depot backorders owed to an arbitrary base, we can find the mean and variance of the base resupply pipeline at base j [we will, of course, be using the mean and variance of depot backorders owed to base j at time $T - OST(T)$, rather than at time T]. We then use the Vari-METRIC approximation to estimate the distributions of base j 's resupply pipeline and of its (base) backorders, just as we have done in the single-base case; we omit the details.

3.6.6 Exponential Repair

The repair times at the base and at the depot as may be treated as random rather than deterministic. An exponential distribution of repair times is a

popular assumption because it leads to tractable computations, both in simulations and in analytical models. Its use does, however, suffer from the drawback that it implies that the most probable repair time is zero. Other distributions exist (such as a two-sided Laplace distribution⁹) that are also amenable to computation but that allow for the most probable repair time to be positive [Crawford].

The model allows only for exponentially distributed *repair* times. (The order and ship process is considerably less variable than would be implied by using an exponential distribution for the order and ship time.) Exponential repair is not consistent with our earlier FIFO assumption, but we will not need that assumption in treating random repair times. We will first describe exponential repair using the continuous approach and then present the model algorithms used in the discrete approach. Then we will expand our algorithms to include exponential repair with suspended resupply and with levels of indenture.

3.6.6.1 CONTINUOUS APPROACH WITH STATIONARY REPAIR TIMES

We use the dynamic version of Palm's theorem [Crawford, Hillestad and Carillo] to compute repair pipeline distributions when the exponentially distributed repair times option is selected. Note that no version of Palm's theorem is needed in the case of time-varying, deterministic resupply times that we have treated thus far; we only require it for random repair times.

Assume that the demands on repair are a nonhomogeneous Poisson process with intensity $\lambda(s)$, that $G(s, t)$ is the probability that an item that fails at time s will be repaired by time t , and that $RpipeRV(t)$ is the number of items in a repair pipeline at time t . *The dynamic Palm's theorem states that $RpipeRV(t)$ is a Poisson random variable with mean (and variance)*

$$\alpha(t) = \int_{-\infty}^t [1 - G(s, t)] \lambda(s) ds. \quad [\text{Eq. 3-54}]$$

To apply Palm's theorem to compute the repair pipeline mean and variance, we must calculate $G(s, t)$ in the case of a **exponential repair time distribution**.

First consider the case where the repair time distribution is the same throughout the scenario, and let $RTRV$ be the repair time (a random variable). In this case we interpret exponential repair to mean that there is a constant β such that the distribution for the repair time is $f(t) = \beta e^{-\beta t}$. Elementary calculus shows that the mean repair time is $E(RTRV) = \frac{1}{\beta}$. We denote the mean repair time $E(RTRV)$ by RT .

⁹This distribution is discussed in Sherbrooke, 1992, in the context of leadtime demand distributions for consumable items.

Now suppose that an item fails at time s . The probability that the item will emerge from repair by time t is just

$$G(s, t) = \Pr(RTRV \leq t - s) = \int_0^{t-s} \beta e^{-\beta x} dx = \int_s^t \beta e^{-\beta(u-s)} du = 1 - e^{-\beta(t-s)},$$

so we have

$$1 - G(s, t) = e^{-\beta(t-s)}, \quad [\text{Eq. 3-55}]$$

and by Palm's theorem, we have

$$E[R_{\text{pipe}}RV(t)] = \text{VAR}[R_{\text{pipe}}RV(t)] = \int_0^t e^{-\beta(t-s)} \lambda(s) ds. \quad [\text{Eq. 3-56}]$$

3.6.6.2 NONSTATIONARY REPAIR TIMES

Non-stationary repair times mean that the repair times form a stochastic process; i.e. the repair time for an item entering repair at time t is a random variable $RTRV(t)$ with probability distribution depending on t .

To treat nonstationary repair times, we use the **hazard function** of a random variable [Parzen]. Let X be a random variable, f be its probability distribution function, and F be its cumulative distribution function. The hazard function H for X is

$$H(x) = \frac{f(x)}{1 - F(x)}. \quad [\text{Eq. 3-57}]$$

The hazard function may be interpreted in terms of conditional probability. To see this, note that times s and t with $s < t$,

$$\int_s^t H(x) dx = \int_s^t \frac{f(x)}{1 - F(x)} dx = -\ln[1 - F(t)] + \ln[1 - F(s)] = \ln \left[\frac{1 - F(s)}{1 - F(t)} \right], \quad [\text{Eq. 3-58}]$$

or

$$\exp \left\{ -\int_s^t H(x) dx \right\} = \frac{1 - F(t)}{1 - F(s)} = \frac{\Pr(X > t)}{\Pr(X > s)} = \Pr(X > t | X > s). \quad [\text{Eq. 3-59}]$$

If X_r is the time to repair an item that enters repair at time $r \leq s$, then Equation 3-59 describes the probability that the item has not emerged from repair by time t given that it is still in repair at time s . If $r = s$, Equation 3-59 becomes

$$\exp \left\{ -\int_s^t H(x) dx \right\} = \Pr(X_s > t) = 1 - G(s, t), \quad [\text{Eq. 3-60}]$$

where $1 - G(s, t)$ is the probability that an item entering repair at s has not emerged by t , as in Equation 3-54. We denote this probability by $UNR(s, t)$, where UNR stands for "un-repair." Equation 3-60 also shows that specifying a hazard function for the repair time X_s uniquely determines the CDF of that repair time.

But the hazard function also has a physical interpretation when the repair time probability distribution f has at most a finite number of discontinuities.¹⁰ Given a time t and a time interval $[t, t + \Delta t]$ short enough that there is at worst a discontinuity at t , but none in the rest of the interval, we have

$$H(t)\Delta t = \frac{f(t)\Delta t}{1 - F(t)} \approx \Pr(t < RTRV \leq (t + \Delta t) | RTRV > t), \quad [\text{Eq. 3-61}]$$

The right-most expression in Equation 3-61 is the probability that an item emerges from repair during $(t, t + \Delta t]$ given that it is still repair at t . Thus $H(t) = \frac{H(t)\Delta t}{\Delta t}$ may be viewed as the average rate of change of that probability over $(t, t + \Delta t]$. Taking the limit as $\Delta t \rightarrow 0$, $H(t)$ is the rate of change of the probability of an item emerging from repair per unit time at time t .

For the exponentially distributed repair times treated in Section 3.6.6.1, we have

$$H(t) = \frac{\beta e^{-\beta t}}{1 - [1 - e^{-\beta t}]} = \beta = \frac{1}{RT}, \quad [\text{Eq. 3-62}]$$

so the hazard function, or the rate of change of the emergence probability for an item in repair, is constant.

Motivated by this observation and the fact that the hazard function of a repair time distribution uniquely determines that distribution (Equation 3-60), we define a **nonstationary exponential repair process** to be a stochastic process $\{RTRV(s)\}_s$ with $RTRV(s) \geq 0$ for all s , and such that for each s , the hazard function H_s for the random variable $RTRV(s)$ is constant (generally not the same constant for each s .) We define $RT(s) = \frac{1}{H_s}$, which would be the mean repair time if the hazard function at s applied for all times.

Substituting Equation 3-60 into Equation 3-54, we find that the mean and variance of the repair pipeline in the case of a nonstationary exponential repair process is

$$\alpha(t) = \int_{-\infty}^t UNR(s, t) \lambda(s) ds = \int_{-\infty}^t \exp \left\{ -\int_s^t \frac{1}{RT(x)} dx \right\} \lambda(s) ds. \quad [\text{Eq. 3-63}]$$

¹⁰Piecewise continuity was not needed until now, but is required for this physical interpretation of the hazard function. Hazard functions in the ASM meet this condition.

Substituting this expression into Equation 3-57 yields the repair pipeline mean and variance.

3.6.6.3 DISCRETE CASE

In the discrete case, the demand intensity $\lambda(t)$ is a step function that remains constant within any given day t . The repair time probability distribution is the same within any given day, changing only from one day to the next, so $RT(t)$ is a step function as well.

To compute the pipeline mean and variance on day t , we must first compute $UNR(s, t)$ for day any day i . For $i-1 \leq s \leq i < t$, Equation 3-60 yields

$$\begin{aligned} UNR(s, t) &= \exp\left(\int_s^t \frac{1}{RT(u)} du\right) \\ &= \exp\left(-\int_s^i \frac{1}{RT(i)} du - \sum_{j=i+1}^t \left[\int_{j-1}^j \frac{1}{RT(j)} du\right]\right) \\ &= \exp\left\{\frac{-(i-s)}{RT(i)}\right\} \cdot \prod_{j=i+1}^t \exp\left\{\frac{-1}{RT(j)}\right\}. \end{aligned} \quad [\text{Eq. 3-64}]$$

Substituting this expression into Equation 3-63, we have

$$\begin{aligned} \alpha(t) &= \int_{-\infty}^t UNR(s, t) \lambda(s) ds \\ &= \sum_{i=-\infty}^t \lambda(i) \int_{i-1}^i UNR(s, t) ds \\ &= \sum_{i=-\infty}^t \lambda(i) \left\{ \int_{i-1}^i \exp\left[\frac{-(i-s)}{RT(i)}\right] \cdot \prod_{j=i+1}^t \exp\left[\frac{-1}{RT(j)}\right] ds \right\} \\ &= \sum_{i=-\infty}^t \lambda(i) \left\{ \prod_{j=i+1}^t \exp\left[\frac{-1}{RT(j)}\right] \int_{i-1}^i \exp\left[\frac{-(i-s)}{RT(i)}\right] ds \right\} \\ &= \sum_{i=-\infty}^t \lambda(i) \left\{ \prod_{j=i+1}^t \exp\left[\frac{-1}{RT(j)}\right] \int_0^1 \exp\left[\frac{-w}{RT(i)}\right] dw \right\}. \end{aligned} \quad [\text{Eq. 3-65}]$$

Let

$$U(i) = \exp\left[\frac{-1}{RT(i)}\right] \text{ and } V(i) = \int_0^1 \exp\left[\frac{-w}{RT(i)}\right] dw = RT(i)[1 - U(i)]. \quad [\text{Eq. 3-66}]$$

Substituting Equations 3-66 into 3-65, we obtain

$$\alpha(t) = \sum_{i=-\infty}^t \lambda(i) \left[\prod_{j=i+1}^t U(j) \right] V(i). \quad [\text{Eq. 3-67}]$$

Since $RT(i)$ and $\lambda(i)$ are known, we have found the mean and variance of the base repair pipeline.

We may interpret $U(j)$ as the probability that a part in repair at the beginning of day j , will not get repaired on day j . Note that U is the "un-repair rate" for a single day, while UNR was the probability an item will not be repaired over the whole time period from s to T . We may think of $V(i)$ as the probability that a unit that fails at a random time on day i will not be repaired before the end of day i . This factor is a minor adjustment the model makes because the actual failure may come any time during that day. If all failures occurred at the beginning of the day, we would add one day to the $U(j)$ product and not need this adjustment.

3.6.6.4 SUSPENDED RESUPPLY AND EXPONENTIAL REPAIR

We now describe how to incorporate exponential repair when resupply is suspended. If we suspend resupply up through day N_s , then U and V jump to 1 during the suspension. Thus in Equation 3-67, we have

$$U(j) = \begin{cases} \exp(-1/RT_p) & \text{when } j \leq 0 \\ 1 & \text{when } 0 < j \leq N_s, \\ \exp(-1/RT_w) & \text{when } j > N_s \end{cases} \quad [\text{Eq. 3-68}]$$

and

$$V(i) = \begin{cases} RT_p\{[1 - U(i)]\} & \text{when } i \leq 0 \\ 1 & \text{when } 0 < i \leq N_s. \\ RT_w\{[1 - U(i)]\} & \text{when } i > N_s \end{cases} \quad [\text{Eq. 3-69}]$$

3.6.6.5 LEVELS OF INDENTURE FOR EXPONENTIAL REPAIR

Exponential repair necessitates an additional calculation when an item contains subassemblies. For instance, to calculate an SRU's pipelines, the model looks back the corresponding LRU's BRT. However, with exponential repair, it is unclear how far to look back, since some LRU repairs take only a day while others take an unlimited number of days. In this case, we use the LRU's effective deterministic base repair time as the "look-back time." This involves computing the LRU's base repair pipeline mean with exponential repair, and then solving for the deterministic repair time that would yield the same pipeline mean.

3.7 SUMMARY

In this chapter, we have derived the probability distribution of base backorders under the assumption that demand rates, resupply times, repair times, and

other inventory system parameters are all changing over time. In the next chapter, we use these distributions to compute aircraft availability under the dynamic conditions encountered in war, as well as those accompanying the introduction of a new aircraft or a force drawdown.

4.0 Availability and Optimization

This chapter discusses the computation of expected availability when cannibalization is permitted, as well as two major extensions to the optimization methodology. We also complete the discussion on the effects of lower indenture items on availability, which began in Chapter 3.

4.1 AVAILABILITY WITH CANNIBALIZATION

In this section we show how to compute expected availability as a function of each item's base and depot spares levels under the assumption that items may be cannibalized. We first motivate the concept of cannibalization by way of an example, and then show how the model's computation of expected availability changes from the method described in Chapter 2 when cannibalization is allowed.

The ability to consolidate broken LRUs onto a single aircraft can improve aircraft availability, albeit at the cost of increasing flight-line maintenance actions. Suppose we have 10 aircraft, each having only 2 different components, and each component has a QPA of 1. One component has two backorders, while the other has one. The backorders for the first component cause holes on two different aircraft, grounding them. But the backorder for the second component may or may not ground a third plane. If one of the planes missing the first component is also missing the second component, then the three holes are on only two planes and the availability is 80 percent. Or the third hole may be on a third plane, yielding a 70 percent availability. If the location of this third hole is random, then there is a 20 percent chance that it is on one of the two already grounded airplanes. Thus the average availability equals 72 percent — the weighted sum of 80 percent availability 20 percent of the time plus 70 percent availability 80 percent of the time. Note that this probability calculation assumes that each component's holes are independent of the others, as in Chapter 2.

Consider the case where two aircraft are NMCS for the first component and one aircraft is NMCS for the second component. Maintenance can restore one aircraft to service by taking a unit of the second component from one of the two planes that are NMCS for the first component and installing it in the aircraft that is down for the second component; this process is called **cannibalization**. Consolidating the holes onto the fewest possible aircraft in this way raises the availability from 72 percent to 80 percent.

Cannibalization consolidates backorders on the fewest possible aircraft, so that the number of aircraft NMCS, at any given time, is the maximum of the

number of aircraft NMCS for each individual component. If $NMCS_l$ is the number of aircraft down for LRU l , then

$$NMCS = \max_l \{NMCS_l\}. \quad [\text{Eq. 4-1}]$$

If we let ENMCS be the expected number of aircraft NMCS, the basic relationship between availability and ENMCS is still valid:

$$\text{availability} = 1 - \frac{ENMCS}{NAC}. \quad [\text{Eq. 4-2}]$$

Unfortunately, the formula for ENMCS is complex, and the simple formulas of Chapter 2 do not apply. We must compute the cumulative distribution function, CDF, of the number of aircraft NMCS and calculate ENMCS from that.

If D is any non-negative integer, the number of aircraft NMCS is less than or equal to D if and only if the number NMCS for each LRU is less than or equal to D . Thus the probability that the number of aircraft down is less than or equal to D is the product, over all the LRUs l , of the probability that the number of aircraft down for that LRU is less than or equal to D :

$$\Pr(NMCS \leq D) = \Pr[\max_l \{NMCS_l\} \leq D] = \prod \Pr(NMCS_l \leq D), \quad [\text{Eq. 4-3}]$$

where the first equality holds by Equation 4-1 and the second inequality follows from our assumption that failures of parts are independent. It remains to find $\Pr(NMCS_l \leq D)$.

Suppose that LRU l has quantity per aircraft QPA_l and base backorders BO_l . The number of slots for LRU l on D aircraft is $D \times QPA_l$, so there will be D or fewer aircraft down precisely when the number of backorders, or holes, is less than or equal to $D \times QPA_l$. Therefore

$$\Pr(NMCS_l \leq D) = \Pr(BO_l \leq D \times QPA_l). \quad [\text{Eq. 4-4}]$$

But from Chapter 3, we know how to compute the distribution of LRU l 's base backorders as a function of the item's base and depot spares levels, so the right-hand side of Equation 4-4 is known.

Substituting this result into Equation 4-3 yields the desired formula for the cumulative distribution of NMCS:

$$\Pr(NMCS \leq D) = \prod_l \Pr(BO_l \leq D \times QPA_l). \quad [\text{Eq. 4-5}]$$

Therefore

$$\begin{aligned} ENMCS &= \sum_{D=1}^{\infty} [D \times \Pr(NMCS = D)] \\ &= \sum_{D=1}^{\infty} \left(D \times \left\{ \prod_l \Pr(BO_l \leq D \times QPA_l) - \prod_l \Pr(BO_l \leq (D-1) \times QPA_l) \right\} \right). \end{aligned} \quad [\text{Eq. 4-6}]$$

Finally, by putting this result for ENMCS into Equation 4-2, *we have found the expected aircraft availability under cannibalization as a function of each item's base and depot spares levels.*

It is convenient to have a more computationally tractable expression for ENMCS. By applying the characterization of an expectation as the sum of the tails of the probability distribution¹

$$\begin{aligned} \text{ENMCS} &= \sum_{D=0}^{\infty} \Pr(\text{NMCS} \geq D) \\ &= \sum_{D=0}^{\infty} [1 - \Pr(\text{NMCS} \leq D)]. \end{aligned} \quad [\text{Eq. 4-7}]$$

We approximate ENMCS with a finite number of terms in the above sum. Taking N large enough that $1 - \Pr(\text{NMCS} \leq D)$ is negligible for $D \geq N$ in Equation 4-7, we have

$$\text{ENMCS} \approx \sum_{D=0}^{N-1} [1 - \Pr(\text{NMCS} \leq D)] = N - \sum_{D=0}^{N-1} \Pr(\text{NMCS} \leq D). \quad [\text{Eq. 4-8}]$$

4.2 COMPUTING SPARES MIXES WITH CANNIBALIZATION

The ASM has three methods of producing spares mixes with cannibalization. These methods correspond to different objective functions — confidence of having a specified number of aircraft available, ENMCS, and a combination of EBO and ENMCS. For any one of these objective functions, the model chooses spares in order of benefit per dollar, where benefit is described in terms of the objective function. The model builds a curve of achieved objective versus total cost, each point on the curve corresponding to a known spares mix. This mix may be thought of as offering either the best performance (highest confidence, lowest ENMCS, etc.) for a given cost, or the lowest cost for a specified level of performance. The model is typically applied in the situation where we seek the spares mix that offers a specified performance for the least cost.

- ◆ Confidence optimization minimizes the cost of meeting a specified probability that the number of aircraft NMCS does not exceed a given target D .² This probability is often called the **confidence** of meeting the target D , and the use of marginal analysis to maximize this confidence is called **confidence optimization**.
- ◆ ENMCS “optimization” seeks to minimize the cost of meeting a target for expected availability. As will be shown, it considers the confidence for all

¹We showed this for backorder CDFs in Section 3.3.4, “Computational Methods.”

²In the USAF context, the associated number of aircraft desired mission capable, $\text{NAC} - D$, is usually called the direct support objective, or DSO.

possible values of NMCS (from zero on up) and “optimizes” a weighted sum of them.

- ◆ EBO/ENMCS “optimization” uses two measures of performance: the primary measure — ENMCS — and a secondary one — EBOs. (This is an experimental method that has not yet been fully tested.)

The methods are similar but result in spares mixes that rely on cannibalization to different degrees. Chapter 5 displays results for the three options (along with results for the no-cannibalization method of Chapter 2).

4.2.1 Confidence Optimization

Given a number of aircraft D that we are willing to allow down, we define the **confidence** of meeting that objective as the probability that there are no more than D aircraft NMCS. From Equation 4-5, we see that

$$\text{confidence} = \Pr(\text{NMCS} \leq D) = \prod_i \Pr(\text{BO}_i \leq D \times \text{QPA}_i). \quad [\text{Eq. 4-9}]$$

Confidence optimization is the original and most basic form of optimization of availability with cannibalization. The formula for confidence under cannibalization is identical in structure to the formula for availability under noncannibalization, and marginal analysis can be applied to obtain optimal spares mixes. (See Appendix A.) Confidence optimization maximizes the probability (confidence) that the number of aircraft NMCS does not exceed a given target.

As we did in Chapter 2 for optimizing aircraft availability without cannibalization, we can optimize the separable function of the logarithm of the confidence:

$$\ln(\text{confidence}) = \ln[\Pr(\text{NMCS} \leq D)] = \sum_i \ln[\Pr(\text{NMCS}_i \leq D)]. \quad [\text{Eq. 4-10}]$$

Let $\Pr(\text{NMCS}_i \leq D | s_i)$ be the probability of no more than D aircraft down for component i , given that its the total spares level, optimally allocated between the depot and base, is s_i . Let ΔX be the change in any quantity X . If we increase the number of units of component k from $s_k - 1$ to s_k , there is no change in $\Pr(\text{NMCS}_i \leq D | s_i)$ except when $i = k$, so we increase the logarithm of the probability of no more than D aircraft down by

$$\begin{aligned} \Delta \ln[\Pr(\text{NMCS} \leq D)] &= \sum_i \ln[\Pr(\text{NMCS}_i \leq D | s_k)] - \sum_i \ln[\Pr(\text{NMCS}_i \leq D | s_k - 1)] \\ &= \ln[\Pr(\text{NMCS}_k \leq D | s_k)] - \ln[\Pr(\text{NMCS}_k \leq D | s_k - 1)] \quad [\text{Eq. 4-11}] \\ &= \Delta \ln[\Pr(\text{NMCS}_k \leq D)]. \end{aligned}$$

Therefore we see that the benefit-to-cost ratio for moving from s_k to $s_k + 1$ units of component k is

$$\frac{\Delta \ln[\Pr(\text{NMCS} \leq D)]}{\Delta \text{cost}} = \frac{\Delta \ln[\Pr(\text{NMCS}_k \leq D)]}{C_k}, \quad [\text{Eq. 4-12}]$$

where C_k denotes the unit cost of component k . This benefit-to-cost ratio is analogous to the sort value in Equation 2-26.

As in the case of aircraft availability without cannibalization, we compute these benefit-to-cost ratios for each spare separately and then sort them to produce an ordered shopping list and confidence-to-cost curve. Also as in that case, we use marginal analysis, suitably modified to handle nonconcave sections of the confidence-to-cost curve, to determine the optimal solution. (See Appendix A for more detail.)

4.2.2 The ENMCS Objective Function

While marginal analysis allows us to produce optimal spares mixes with respect to confidence, confidence has some drawbacks as an objective function. Note in Equations 4-9 and 4-12 that confidence and the associated marginal benefit depend only on the probability that the number of aircraft NMCS is less than the target number of aircraft allowed down. Confidence optimization for an NMCS target of four aircraft down does not consider the individual probabilities of zero, one, two, etc., NMCS aircraft down but only the sum of those probabilities. A wing commander, however, may well be interested in how likely the wing is to have a particular number of aircraft available. Furthermore, even if the allowed number of aircraft down is exceeded, it makes a difference operationally if it is exceeded by only 1 aircraft or by 10 aircraft, and this issue is not directly addressed by confidence optimization.

This shortcoming of confidence optimization leads to consideration of ENMCS as an objective function, since ENMCS is more sensitive than confidence-level to the distribution of NMCS. Unfortunately, ENMCS is an inseparable objective function, as we can see from Equation 4-6.

We have experimented with a “greedy heuristic” that *searches* for the spares mix that attempts to minimize ENMCS in the same way as marginal analysis. This method produces good solutions, though not demonstrably optimal ones. The term “greedy heuristic”³ emphasizes the fact that the search algorithm “greedily grabs” a unit of the component with the best benefit-to-cost ratio at each stage of the search; *we reserve the term “marginal analysis” for the special case where the search may be shown to locate the optimal solution.*

³In some contexts, this is called the “method of steepest descent” or “gradient search.”

As in the case of confidence optimization, we develop an expression for the benefit-to-cost ratio, or sort value, for adding an additional unit of component k . In Equation 4-8, set $N = NAC + 1$, where NAC is the number of aircraft; then

$$ENMCS = NAC + 1 - \sum_{D=0}^{NAC} \Pr(NMCS \leq D). \quad [\text{Eq. 4-13}]$$

Thus minimizing $ENMCS$ is the same as maximizing the sum of $p(NMCS \leq D)$. Using the fact that $E(\text{mission-capable aircraft}) = NAC - ENMCS$, we can rewrite Equation 4-13 as:

$$E(\text{mission-capable aircraft}) = -1 + \sum_{D=0}^{NAC} \Pr(NMCS \leq D). \quad [\text{Eq. 4-14}]$$

Our objective from here on will be to maximize this sum. Using Equation 4-3, we rewrite Equation 4-14 as

$$\begin{aligned} E(\text{mission-capable aircraft}) &= -1 + \sum_{D=0}^{NAC} \Pr(NMCS \leq D) \\ &= -1 + \sum_{D=0}^{NAC} \left[\prod_l \Pr(NMCS_l \leq D) \right] \\ &= -1 + \sum_{D=0}^{NAC} \exp \left\{ \ln \left[\prod_l \Pr(NMCS_l \leq D) \right] \right\} \\ &= -1 + \sum_{D=0}^{NAC} \exp \left\{ \sum_l \ln \left[\Pr(NMCS_l \leq D) \right] \right\}. \end{aligned} \quad [\text{Eq. 4-15}]$$

Let s_j be the spares level for LRU j , let B_j be the increase in $E(\text{mission-capable aircraft})$ due to an additional spare of LRU j , and let

$$\begin{aligned} &\Delta \exp \left\{ \sum_l \ln [\Pr(NMCS_l \leq D)] \right\} \\ &= \exp \left\{ \sum_l \ln [\Pr(NMCS_l \leq D | s_j + 1)] \right\} - \exp \left\{ \sum_l \ln [\Pr(NMCS_l \leq D | s_j)] \right\}. \end{aligned} \quad [\text{Eq. 4-16}]$$

Then from Equations 4-15 and 4-16, we have

$$B_j = \sum_{D=0}^{NAC} \Delta \exp \left\{ \sum_l \ln [\Pr(NMCS_l \leq D)] \right\}. \quad [\text{Eq. 4-17}]$$

Using the differential relationship $\Delta e^{f(x)} \approx \frac{d}{dx}(e^{f(x)})\Delta x = e^{f(x)}f'(x)\Delta x \approx e^{f(x)}\Delta f(x)$, we can approximate this increase by

$$\begin{aligned} B_j &\approx \sum_{D=0}^{NAC} \exp \left\{ \sum_l \ln [\Pr(NMCS_l \leq D)] \right\} \times \Delta \sum_l \ln [\Pr(NMCS_l \leq D)] \\ &= \sum_{D=0}^{NAC} \exp \left\{ \sum_l \ln [\Pr(NMCS_l \leq D)] \right\} \times \Delta \ln [\Pr(NMCS_j \leq D)] \quad [\text{Eq. 4-18}] \\ &= \sum_{D=0}^{NAC} \Pr(NMCS \leq D) \times \Delta \ln [\Pr(NMCS_j \leq D)], \end{aligned}$$

where the first strict equality follows because the only term in the sum $\sum_l \ln [\Pr(NMCS_l \leq D)]$ that changes when the spares level for component j changes from s_j to $s_j + 1$ is the one where $l = j$, and the second strict equality follows by Equation 4-10.

We may write Equation 4-18 as

$$B_j \approx \sum_{D=0}^{NAC} W_D \times \Delta \ln [\Pr(NMCS_j \leq D)], \quad [\text{Eq. 4-19}]$$

where

$$W_D = \Pr(NMCS \leq D). \quad [\text{Eq. 4-20}]$$

That is, the benefit function for adding a unit of component j is a weighted average of the benefit functions for all possible NMCS values D , and the weights are the corresponding confidence levels (the cumulative probability that D or fewer aircraft are down for lack of spares). Since we compute these probabilities to compute ENMCS, they are readily available to use as weights. Another way to describe the benefit function is as a dot product of two vectors; the weights vector and the delta log confidence vector (which we call the benefit vector).

Unfortunately, each additional buy changes the weights vector, and thus with each buy all the dot products must be recomputed. As the model builds the curve of ENMCS versus cost, buying spares one at a time, it must recompute all the dot products after each buy. This procedure is prohibitively slow for large numbers of components.⁴

Note that in adding spares for components other than component j to the mix, leaving the number of units of component j unchanged, the benefit vector does not change, but the weight vector does. The change is a consequence of the inseparability of the objective function. If the objective function could be written

⁴Despite this drawback, the ASM does use this approach to compute AWP time for SRUs. The relationship for LRUs AWP for SRUs is the same as the formula for aircraft NMCS for LRUs. Since there are normally few SRUs on each LRU, we use the greedy heuristic to produce the curve of AWP versus cost for each LRU, as we will describe later in this chapter. Because of the small number of SRUs, the inefficiency of the greedy heuristic is not significant.

as a sum (or product) of functions, each depending only on the spares level of a single LRU, then the marginal benefits could also be computed independently, as in the case of confidence-level optimization.

We note that we can also express the marginal benefits for confidence-level optimization in the dot product form. Set the weight vector to $\delta(D)$, the Dirac delta, with the coordinate equal to one in the D th position and zero otherwise; then the benefit from an additional spare shown in Equation 4-19 reduces to the change in the log of the confidence level.

Thus Equation 4-19 (with the appropriate choice of weights) describes the marginal benefit of adding a spare under both confidence-level optimization and the greedy heuristic. In Sections 4.2.3 and 4.2.4, we will discuss the behavior of benefit functions with other weights.

4.2.3 Weighted Heuristics

Between the two extremes of the greedy heuristic, with its nonconstant weight vector, and confidence-level optimization, with its simple Dirac delta-weight vector, lies a family of solution algorithms determined by different weight vectors. Our investigation into these weighted heuristics revealed two types of particular interest.

One type of heuristic seeks to minimize ENMCS with less computational effort than the method we described previously. The other heuristic allows us to produce spares mixes resulting in a better balance of confidence-level, ENMCS, and EBOs, for a given investment in spares, than we obtain with confidence or ENMCS optimization.

In either case, we limit ourselves to fixed-weight vectors W_D for processing efficiency, and perform a greedy algorithm, choosing spares in order of their benefit-to-cost ratios, where the ratios for any component j are of the form

$$\frac{\sum_{D=0}^N W_D \times \Delta \ln[\Pr(NMCS_j \leq D)]}{cost_j}. \quad [\text{Eq. 4-21}]$$

Viewing the numerator as a dot product, as we did earlier, we see that the benefit vector (i.e., the log confidence vector) for component j is a function only of component j 's stock level, and need be computed only once for each spares mix that has a given number of units of component j .

This last fact ensures efficient processing and short run times, although we cannot guarantee strict optimality of the solutions. The key to the goodness of the solutions is a judicious choice of weights, which we discuss in Sections 4.2.3.1 and 4.2.3.2.

4.2.3.1 WEIGHTED HEURISTIC WITH AN ENMCS OBJECTIVE

To perform ENMCS “optimization,” the model uses a weighted heuristic with a set of weights formulated for the user’s ENMCS target. For each ENMCS target in the typical range of interest, there are weights that yield spares mixes with approximately the same cost as those produced by the slower algorithm described in Section 4.2.2 (we will refer to the method of that section as the “pure” greedy heuristic in the discussion that follows).

One way to generate an effective set of weights is as follows: start with any initial weight vector, run the weighted heuristic search until attaining the ENMCS target, store the confidence vector (the CDF) for this spares mix, and then rerun the model (to the same ENMCS target) using that stored CDF as the input weight vector. We have observed that after this process of computing CDFs has been iterated several times (for any of the model databases we have tried), the sequence of weight vectors changes little with subsequent iterations, and any of the CDFs obtained after this point produce spares mixes meeting the ENMCS target with comparable cost to that obtained by the pure greedy heuristic.

We experimented with various databases and NMCS targets and developed an empirical method for generating weights that closely matches the more cumbersome iterative process but is based solely on the NMCS target. (We will present examples of those weights shortly.) Runs using those weights produced spares mixes that met the ENMCS target with cost close to that obtained by the pure greedy heuristic and had a significantly lower ENMCS than mixes with the same cost produced using confidence optimization. This technique also has the unintended benefit of yielding fewer total expected backorders than does confidence optimization. Since fewer expected backorders means fewer expected cannibalization actions (all else being equal), this benefit is of considerable interest.

Why does the ENMCS weighted heuristic yield a lower number of EBOs than does confidence optimization? Recall that for *confidence* optimization, we focus on only one element in the benefit vector. All weights are set to zero except the one for the target (assuming an integer target for simplicity).

We found that the weights for the low NMCS values were the key to driving expected backorders down for the ENMCS weighted heuristic. A nonzero W_0 or W_1 in Equation 4-21 has a great impact on expected backorders. Many parts are unlikely to have more than a small number of units in their pipelines and thus have almost no probability of grounding more than 1 aircraft. In confidence optimization with an NMCS target of, say, 4, there is no incentive for the algorithm to buy spares of these parts. However, in the ENMCS heuristic, the ASM considers the probability of grounding 1, 2, or 3 aircraft. Items unlikely to have more than a small number in their pipelines contribute to these probabilities, and the model will tend to add them to the spares mix. Thus the ENMCS heuristic produces an increased range of parts with resulting lower total expected backorders. Chapter 5 presents model run examples that quantify the change in

expected backorders that we typically observe in moving from confidence optimization to the ENMCS weighted heuristic.

The empirically derived weight formulas are fairly complex, so we will not present them here. To give the reader an idea of the typical weights, we present an example. The ENMCS heuristic with a target NMCS of 4 uses the set of weights in Table 4-1. The i th component of this vector approximates the $CDF(i)$ of NMCS resulting from a spares mix produced by the greedy algorithm with an ENMCS of 4. The vector is similar for other relatively small ENMCS targets, except that it is shifted so that the 0.73 weight corresponds to the ENMCS target. However, ENMCS targets much higher than 4 create slightly broader CDFs, as will be shown in Section 4.2.3.2.

Table 4-1.
*Comparison of Weights for Confidence
and ENMCS Methods (NMCS Target = 4)*

Number of aircraft NMCS, D	ENMCS method weights, W_D	Confidence method weights, W_D
0	0.00	0
1	0.0000005	0
2	0.05	0
3	0.40	0
4	0.73	1
5	0.90	0
6	0.95	0
7	0.98	0
8	0.99	0
9	1.00	0
10	1.00	0
$\vdots$	$\vdots$	$\vdots$

4.2.3.2 WEIGHTED HEURISTIC WITH JOINT EBO/ENMCS OBJECTIVE

We also developed a weighted greedy heuristic that focuses on two measures of performance: the primary measure — ENMCS — and a secondary one — EBOs. EBOs are important because they are a surrogate for maintenance workload. Although cannibalization yields improved aircraft availability over the similar noncannibalization case, it has the disadvantage of causing flight-line maintenance actions to increase. Cannibalization requires the additional maintenance steps (and risks) of borrowing items from grounded aircraft and then reinstalling them later when the backorder is filled. Other things being equal, a

spares mix with a high projected expected backorder total requires more cannibalization actions to produce a given availability than does a mix with lower expected backorders. With the EBO/ENMCS heuristic, the model develops a compromise spares mix that produces a slightly worse ENMCS output than does the pure ENMCS heuristic — but significantly lower expected backorders (i.e., a reduced maintenance workload).

To develop the weights for the EBO/ENMCS heuristic, we enhance the EBO reduction effect in the ENMCS heuristic by increasing the weights for the lower NMCS states. We achieved good results by increasing the lowest positive weight and attaching an exponential tail to the left of the weight vector. That is, we take the lowest nonzero weight from the ENMCS optimization weights, set its weight to 0.005, make the previous weight 70 percent of that, and each of the other previous weights 70 percent of its successor. Thus for an NMCS target of 4, we start with the NMCS weights presented in Table 4-1 and change the weights for $D = 0, 1, \dots$, as shown in Table 4-2.

Table 4-2.
Weights for ENMCS and EBO/ENMCS Methods
(NMCS Target = 4)

Number of aircraft NMCS, D	ENMCS method weights, W_D	EBO/ENMCS method weights, W_D
0	0.00	0.0035
1	0.0000005	0.005
2	0.05	0.05
3	0.40	0.40
4	0.73	0.73
5	0.90	0.90
6	0.95	0.95
7	0.98	0.98
8	0.99	0.99
9	1.00	1.00
10	1.00	1.00
$\vdots$	$\vdots$	$\vdots$

For an ENMCS target of 7, the weights for the pure ENMCS and the EBO/ENMCS heuristics are as shown in Table 4-3.

Table 4-3.
Weights for ENMCS and EBO/ENMCS Methods
 (NMCS Target = 7)

Number of aircraft NMCS, D	ENMCS method weights, W_D	EBO/ENMCS method weights, W_D
0	0.00	0.0012005
1	0.00	0.001715
2	0.00	0.00245
3	0.00	0.0035
4	0.001	0.005
5	0.080	0.080
6	0.400	0.400
7	0.715	0.715
8	0.885	0.885
9	0.950	0.950
10	0.980	0.980
11	0.990	0.990
12	1.00	1.00
13	1.00	1.00
$\vdots$	$\vdots$	$\vdots$

Changing the weights from those of the pure ENMCS case to those shown for the EBO/ENMCS case drops expected backorders precipitously, with only a slight ENMCS penalty. (See Chapter 5 for examples with detailed expected backorder comparisons.)

Note that the ENMCS weights are somewhat different from those with a target of 4, shown in Table 4-1. This subtle difference is important — for higher NMCS targets, we have found that we need weights with a slightly broader distribution. Failure to allow for this broader distribution can significantly degrade performance for high NMCS targets.

4.2.4 Confidence Optimization with Nonintegral Targets

Another application of weights is confidence optimization with nonintegral targets. In this case, all weights are set to zero except the ones for NMCS values immediately below and above the target. The nonzero weights are set so they sum to one, and the weighted average of their corresponding NMCS values equals the NMCS target. For example, a target of 3.6 gives the weights shown in Table 4-4.

Table 4-4.
Weights for Confidence Method
(NMCS Target = 3.6)

Number of aircraft NMCS, D	Confidence method weights, W_D
0	0.0
1	0.0
2	0.0
3	0.4
4	0.6
5	0.0
6	0.0
$\vdots$	$\vdots$

4.2.5 Partial Cannibalization

In Chapter 2 we introduced the concept of system availability and discussed the calculation of availability without cannibalization. We have now extended the discussion to include cannibalization, whereby maintenance can consolidate holes onto as few aircraft as possible. Some components, however, may be difficult or impossible to cannibalize or may involve such a risk of collateral damage as to make cannibalization unwise. The model has a partial cannibalization option, which is still experimental; we explain the underlying ideas below.

In this case, we must compute ENMCS somewhat differently from the way we did in the second part of Equation 4-8. But as in that full cannibalization case, we still need to compute $\Pr(\text{NMCS} \leq D)$ for each $D \geq 0$.

Assume, as in the full cannibalization case, that each part has an application percentage of 100, and again assume that failures of different parts are independent. Assume that there are NAC aircraft, there are QPA_l locations for part l per aircraft, and part l has BO_l backorders. We will further assume that for each l , BO_l is small relative to $NAC \times QPA_l$, for each noncannibalizable part, $BO_l < NAC$, and that backorders for noncannibalizable parts are uniformly distributed over aircraft.

Let $NMCS_{NC}$ be the number of aircraft with at least one hole for a noncannibalizable part and $NMCS_C$ be the number of aircraft with at least one hole for a cannibalizable part, after consolidating holes onto the minimum number of aircraft. Let $NMCS$ be the minimum number of aircraft down; we will derive an expression for $\Pr(\text{NMCS} \leq D)$.

If $NMCS_C \leq NMCS_{NC}$, then as in the full cannibalization case, all of the holes for cannibalizable parts can be moved (if they are not already there) to aircraft down for noncannibalizable parts, so that in this case $NMCS = NMCS_{NC}$. On the other hand, if $NMCS_C > NMCS_{NC}$, then by moving as many holes as possible to the aircraft down for noncannibalizable parts, we have $NMCS_{NC}$ aircraft down for both noncannibalizable and cannibalizable parts, and $NMCS_C - NMCS_{NC}$ aircraft down for cannibalizable parts only. Therefore in this case, $NMCS = NMCS_{NC} + (NMCS_C - NMCS_{NC}) = NMCS_C$. Combining the two cases shows that

$$NMCS = \max(NMCS_{NC}, NMCS_C). \quad [\text{Eq. 4-22}]$$

From Equation 4-22 and the assumption of the independence of failures for different parts,

$$\begin{aligned} \Pr(NMCS \leq D) &= \Pr(NMCS_{NC} \leq D, NMCS_C \leq D) \\ &= \Pr(NMCS_{NC} \leq D) \Pr(NMCS_C \leq D). \end{aligned} \quad [\text{Eq. 4-23}]$$

As in the development of Equation 4-5 in the full cannibalization case, we see that the second factor of the product in Equation 4-23 is given by

$$\Pr(NMCS_C \leq D) = \prod_l \Pr(BO_l \leq D \times QPA_l), \quad [\text{Eq. 4-24}]$$

where now the index l refers only to cannibalizable parts.

We approximate the first factor in Equation 4-23.⁵ Whenever we refer to a part l in this discussion, it means a noncannibalizable part. Let F be the fraction of the NAC aircraft *not* down for a part. Because backorders are uniformly distributed over aircraft and $BO_l < NAC$, the fraction of aircraft not down for part l is $F_l = 1 - \frac{BO_l}{NAC}$; since BO_l is a random variable, this quantity is as well. Because of our assumption that backorders for different parts are independent and uniformly distributed across aircraft, the fraction of the F_l NAC aircraft not down for part l that is not down for another part k , is F_k , and the fraction of the total aircraft not down for either part is $F_l F_k$. Thus

$$F = \prod_l \left(1 - \frac{BO_l}{NAC} \right). \quad [\text{Eq. 4-25}]$$

⁵ For a treatment not assuming uniformly distributed backorders, see Gaver.

To approximate F , observe that for x close enough to 1, the first two terms of the Taylor-series expansion for $\ln(x)$ about $x = 1$ form a reasonable approximation; thus $\ln(x) \approx x - 1$. In Equation 4-25, let $x = 1 - \frac{BO_i}{NAC}$. Then we have

$$1 - \frac{BO_i}{NAC} = x = e^{\ln(x)} \approx e^{(x-1)} = e^{-\frac{BO_i}{NAC}}. \quad [\text{Eq. 4-26}]$$

Applying Equation 4-26 to each factor in Equation 4-25, we find

$$F = \prod_i \left(1 - \frac{BO_i}{NAC}\right) \approx \prod_i e^{-\frac{BO_i}{NAC}} = e^{-\frac{\sum BO_i}{NAC}}. \quad [\text{Eq. 4-27}]$$

Letting $BO = \sum_i BO_i$ in Equation 4-27, we find

$$NMCS_{NC} = NAC \times (1 - F) \approx NAC \times \left(1 - e^{-\frac{BO}{NAC}}\right). \quad [\text{Eq. 4-28}]$$

Since BO is a sum of independent random variables, uniformly bounded by NAC , we may approximate its probability distribution by a normal distribution with mean EBO and variance $VAR(BO)$.⁶ We compute the means and variances of each item's backorders as in Chapter 3, and using the independence of backorders for different components, sum the means and variances respectively, to obtain EBO and $VAR(BO)$. By Equation 4-28 and some straightforward algebra,

$$\Pr(NMCS_{NC} \leq D) = \Pr\left[BO \leq -NAC \ln\left(1 - \frac{D}{NAC}\right)\right]. \quad [\text{Eq. 4-29}]$$

Substituting the probabilities from Equations 4-29 and 4-24 into 4-23, we have found $\Pr(NMCS \leq D)$.

The current version of the model uses a somewhat simpler method, in which we compute the **expected** fractions of aircraft not down for each part, replace the number of backorders for each part in the above derivation by its expected value, and obtain

$$ENMCS_{NC} = NAC \left(1 - e^{-\frac{EBO}{NAC}}\right). \quad [\text{Eq. 4-30}]$$

The model then approximates $\Pr(NMCS_{NC} \leq D)$ with a cumulative Poisson distribution with mean $ENMCS_{NC}$. Substituting this probability and the result of Equation 4-24 into Equation 4-23, we obtain $\Pr(NMCS) \leq D$.

Whether using the normal or Poisson approximation for $\Pr(NMCS \leq D)$, the backorder distributions from which these probabilities are computed are functions of the spares levels, so the confidence of no more than D aircraft down is

⁶Provided that the number of noncannibalizable parts is sufficiently large, this is justified by a generalization of the central limit theorem due to Lindeberg [Feller, Vol. 1].

now expressed as a function of those spares levels. For confidence optimization, this is sufficient; for ENMCS optimization, we use the ENMCS computed by Equation 4-13. We recommend ENMCS optimization for the reasons discussed in the pure cannibalization case.

To illustrate partial cannibalization in a simple case, we again consider the example from Section 4.1 with 10 aircraft with only 2 different parts. One component has two backorders, while the other has one. With cannibalization, only 2 aircraft need be NMCS; without it, most likely 3.

Suppose only one of the components may be cannibalized. If that component is the one with one backorder, the serviceable unit needed may be moved from 1 of the 2 aircraft NMCS for the other component to restore the third aircraft to service. Conversely, if the component with 2 backorders may be cannibalized, one of the aircraft NMCS for that component may be restored to service by cannibalizing from the single aircraft NMCS for the other component. In either case, the total NMCS is 2 — the same as for full cannibalization.

In fact, in any case in which only 1 component may not be cannibalized, the holes for all the other parts can be consolidated as much as possible to conform to the NMCS for that 1 component, and the NMCS number is the same as for full cannibalization.

4.3 LEVELS OF INDENTURE

Most aircraft repair is modular. The aircraft is repaired on the flight line by removing and replacing failed LRUs. LRUs are repaired in the shop by removing and replacing failed SRUs. This relationship is called levels of indenture, because work breakdown structure depictions of parts on an aircraft typically indent subassemblies beneath their main assemblies in the listings. It is not unusual for a system to have more than two levels of indenture. SRUs may contain sub-SRUs, which may contain sub-sub-SRUs, and so on. We will not use any special names for these **lower indenture** items, simply calling them SRUs.

4.3.1 Awaiting Parts Pipelines for LRUs

Since only LRUs are used directly in aircraft repair, only LRU backorders directly make aircraft NMCS. SRU backorders hold up the repair of LRUs but do not directly affect aircraft and thus should *not* be included in the formulas for aircraft availability, except as their backorders affect backorders of their parent LRUs. The equation for aircraft availability in the noncannibalization case is unchanged, but we must now understand that the subscript L stands for LRU_L and that the product is taken over the LRUs only:

$$A = \Pi \left(1 - \frac{EBO_L}{TI_L} \right)^{QPA_L} \quad [\text{Eq. 4-31}]$$

Similarly, in the cannibalization case, Equations 4-1 through 4-30 are unchanged, but again we understand that the L in them stands for LRU. Thus for example, in Equation 4-3, the product of $\Pr(NMCS_l \leq D)$ is taken over the LRUs only.

The model explicitly accounts for the delay in LRU repair caused by SRU backorders. SRU backorders cause holes in LRUs in much the same manner that LRU backorders cause holes in aircraft. Those LRUs with holes for SRUs (awaiting parts, or AWP) constitute another segment of the LRU pipeline. We assume that this additional pipeline segment pertains to base repair only; depot repair is assumed to have a sufficient supply of SRUs. This permits us to avoid the complication of a multi-echelon AWP pipeline calculation.

In Chapter 3, we computed the distributions of SRU backorders. From these distributions, we compute the expected number of LRUs AWP from the SRU backorders (we explain this below), add that pipeline to the others, and calculate LRU backorder distributions and aircraft availability as before. (This assumes the model option selected is a VMR of 1. Otherwise, the model adds the LRU's AWP pipeline variance to the variance of the base resupply pipeline and calculates the distribution of LRU backorders as in Chapter 3. For simplicity, we will not treat the Vari-METRIC option in the remainder of the chapter.)

MOD-METRIC [Muckstadt] and the Aircraft Availability Model (AAM) [O'Malley] assume that each SRU backorder holds up exactly one LRU in base repair. A particular LRU's AWP pipeline is the sum of all of *its* SRUs' expected backorders (SRU backorders for a *different* LRU do not affect *this* LRU). Thus if we let AWP_L denote the mean awaiting parts pipeline for LRU L , $TEBO_s$ denote the number of expected backorders for SRU_s , and $S \subset L$ denote an SRU that is a sub-component of LRU_L , we have

$$AWP_L = \sum_{S \subset L} TEBO_s. \quad [\text{Eq. 4-32}]$$

The ASM, in contrast, assumes that base repair consolidates SRU shortages (i.e., cannibalizes SRUs). This process is analogous to LRU cannibalization on aircraft, and the formula for the expected number of an LRU AWP is analogous to the cannibalization ENMCS formula.

To see this, let $NAWP_L$ be the number of units of LRU L AWP, $NAWP_{L,s}$ be the number of units of LRU L AWP for SRU S , and $QPA_{L,s}$ be the quantity of SRU S installed on LRU L ; then we have the analog of Equations 4-3 and 4-5:

$$\begin{aligned} \Pr(NAWP_L \leq D) &= \prod_{S \subset L} \Pr(NAWP_{L,s} \leq D) \\ &= \prod_{S \subset L} \Pr(BO_s \leq D \times QPA_{L,s}). \end{aligned} \quad [\text{Eq. 4-33}]$$

Just as we derived the expression for ENMCS in Equation 4-8, we use these probabilities to get AWP_L , the expected number of units of LRU L AWP:

$$AWP_L = \sum_{D \geq 0} [1 - \Pr(NAWP_L \leq D)]. \quad [\text{Eq. 4-34}]$$

Note that the probabilities in Equation 4-34, and therefore AWP_L , depend on the spares levels for the SRUs belonging to LRU L .

Including the AWP pipeline, and assuming that there are N uniform bases,⁷ Equation 2-12 for the mean base resupply pipeline at base j becomes

$$Bpipe_{L,j} = BRpipe_L + OSpipe_L + \frac{DEBO_L}{N} + AWP_L, \quad [\text{Eq. 4-35}]$$

where the subscript L has been included in all the variables to reference LRU_L .

As in Chapter 3, this expected pipeline is used to compute the distribution of base backorders; the only difference is that the distribution now depends on the spares levels for the LRU's associated SRUs as well as the base and depot spares levels for the LRU itself.

4.3.2 Multi-Indenture Optimization

In Section 2.2.3, considering only LRUs without SRUs, the ASM finds the combination of base and depot spares levels with the lowest EBO for each total LRU spares level; this step is a prelude to marginal analysis. The ASM trades off depot spares against base spares. For an LRU with subassemblies, the model performs a three-way tradeoff between the depot, base, and SRUs' spares levels to find the combination that gives the lowest total LRU base EBOs *for a specified dollar total*. For example, for an LRU that costs \$1,000, the model trades off spares at the depot versus spares at the base versus \$1,000 increments of SRUs.

The ASM first processes SRUs, performing the depot/base tradeoff and marginal analysis for each SRU. Then it aggregates the SRU results for each particular LRU, producing an AWP-versus-cost curve for that LRU along with its ordered shopping list of SRU buys. The ASM uses this curve of AWP-versus-cost, together with its knowledge of how base and depot spares levels affect the LRU's expected backorders, in the LRU's three-way tradeoff. If there are more than two levels of indenture, the model first processes the lowest level SRUs, performing the depot/base tradeoff and marginal analysis. The results are then passed up to the next level for input into the three-way tradeoffs. These tradeoffs and marginal analysis are performed at this level and the results are passed up to the next level. This process is repeated until level 1 (the LRU level) is reached.⁸

⁷See Chapter 3 for computing the expected base resupply pipeline with multiple bases.

⁸See O'Malley for more detail.

4.3.3 Common-Component SRUs

The ASM also allows for **common-component** SRUs (SRUs used in more than one LRU). The model computes the expected pipelines for these SRUs, considering demands generated by the repair of each of their parent LRUs and allocating backorders to those parents.

In MOD-METRIC and the AAM (which do not consider cannibalization of SRUs), computing AWP for common SRUs is straightforward. SRU backorders are simply prorated to LRU "parents," using factors computed based on usage rates. For example, suppose that 60 percent of SRU 1 total utilization is on LRU A and 40 percent on LRU B (typically this means that 60 percent of the flying time comes on LRU A and 40 percent on LRU B, although use can also relate to other factors besides flying hours). We assume that 60 percent of the demands for SRU 1 came from LRU A and that 60 percent of the backorders have an impact on LRU A. Thus for SRU 1, we give LRU A a 60 percent prorating factor and LRU B a 40 percent one. We then apply 60 percent of SRU 1's backorders to LRU A and 40 percent to LRU B. Each prorated piece of the SRU's EBOs is then included in Equation 4-35.

However, because the ASM considers SRU cannibalization, the formula is much more complicated. Instead of simply prorating the SRU EBOs, the model must prorate the entire CDF of SRU backorders (since the AWP calculation in Equations 4-33 and 4-34 requires that CDF). To do this, we have to break down the CDF into the probability distribution function, split the probability distribution binomially, and then roll up the results into CDFs again. Given a particular LRU with an SRU, this split divides the backorders for the SRU into those that affect the given LRU and those that affect other LRUs. (The ASM does make the simplifying assumption that cross-cannibalization — the cannibalization of an SRU from LRU type A to an LRU of a different type — does not occur.)

To make this more precise, let $P_{L,S}$ be the prorating factor (just described) for SRU_s on LRU_L , BO_s be the total backorders for SRU_s , and $BO_{L,S}$ be the backorders of SRU_s that affect LRU_L . Viewing $P_{L,S}$ as the probability that a randomly chosen backorder for SRU_s affects LRU_L , we see that

$$\begin{aligned} \Pr(BO_{L,S} = N) &= \sum_{M=N}^{\infty} \Pr(BO_{L,S} = N | BO_s = M) \Pr(BO_s = M) \\ &= \sum_{M=N}^{\infty} \binom{M}{N} P_{L,S}^N (1 - P_{L,S})^{M-N} p(BO_s = M). \end{aligned} \quad [\text{Eq. 4-36}]$$

We then substitute the probability distribution of the SRU_s backorders affecting LRU_L into Equation 4-33, and use it in computing the LRU_L 's expected AWP pipeline (Equation 4-34).

5.0 Sample Results and Sensitivity

In previous chapters, we have described the mathematics of the ASM and several modes of operating the model. We will now present and analyze some representative spares mixes generated by the model. We will describe and present examples of the basic model inputs and how the spares mix varies with those inputs. We will discuss the following:

- ◆ Basic model inputs
 - ▶ The availability and cost target options
 - ▶ The user-supplied parameters necessary for specifying a spares mix
- ◆ Developing spares mixes for steady-state conditions, dynamic wartime conditions, and combinations of the two
- ◆ Optimization with multiple levels of indenture
- ◆ Spares mixes with and without cannibalization
- ◆ Dynamic flying programs over time
- ◆ Spares mixes supporting peacetime and wartime activity simultaneously.

5.1 MODEL INPUTS

5.1.1 Basic Targets

The ASM computes spares requirements for two types of targets: either an unconditional availability target or an availability target under a budget constraint. An availability target is expressed in terms of NMCS aircraft or aircraft on ground (AOG). Translation from availability to an NMCS value is based on equations 2-15, 2-16, and 4-2, which state that:

$$availability = 1 - \frac{NMCS}{NAC}, \quad [Eq. 5-1]$$

or

$$expected\ availability = 1 - \frac{ENMCS}{NAC}, \quad [Eq. 5-2]$$

where NAC equals the number of aircraft under consideration. The model always attempts to meet the availability target or targets for the least cost. With an unconstrained budget, it always meets the availability targets.

In addition to expressing the availability target as an average, the user can specify a confidence for that target. For example, the user may want to be 95 percent certain that no more than four aircraft in a squadron will be NMCS over the course of a war scenario. Unfortunately, the confidence level and the related availability rate are frequently confused with each other. For example, suppose we have a 20-aircraft squadron with a desired availability of 16 out of 20 aircraft, or 80 percent of the squadron. The statement of the goal, then, would be "95 percent certain of 80 percent availability." The potential for confusion is obvious.

5.1.2 Flying-Hour Scenarios

Since component failures are driven by weapon-system activity levels, the ASM requires that a flying-hour scenario be specified. The model can estimate spares requirements for the following:

- ◆ Steady-state conditions under which the aircraft flying hours (and item pipelines) remain constant over time.
- ◆ Dynamic conditions under which aircraft flying hours (and item pipelines) may vary dramatically over a few days, weeks, or months.
- ◆ Combined scenarios under which aircraft flying hours change from steady-state to dynamic conditions. This allows the user to estimate the effect of the steady-state conditions on aircraft availability during the period of the dynamic conditions and to compute the spares required for both periods.

Figure 5-1 illustrates flying-hour profiles for steady-state and a 30-day war for three different weapon systems (aircraft) — A, B, and C.

The flying profile for aircraft A has a steady state period followed by a dynamic one. This scenario is typical and is usually interpreted as a period of peacetime followed by the outbreak of a conflict. Our convention is that the steady-state period covers all the days preceding and including day zero, and that the dynamic period starts on day 1. The dynamic period can last up to 99 days, with a different flying-hour program for each day if necessary.




Figure 5-1.
Sample Flying-Hour Profile for Aircraft A, B, and C

The ASM can generate and analyze requirements for one or two days specified by the user. While it is the availabilities on these specified analysis days that explicitly drive the requirements calculation, the model considers activity over the entire period, since this activity contributes to the pipelines on the analysis days. For a one-day analysis, the ASM computes the spares requirements to meet a user-specified availability or budget target on that day.

For a two-day analysis, the ASM requires availability or budget targets for each analysis day. The model first purchases enough spares to meet the target on the first analysis day or — if constrained by the budget — purchases the spares that yield the best performance without exceeding the budget constraint. For the second analysis day, the model includes those spares already computed in support of the first day in the inventory, and it also considers any changes the maintenance system may undergo between the two days (e.g., a transition from a no-cannibalization mode to a “cannibalize-when-necessary” mode). It then computes the additional spares required to meet the specified second-day target or second-day incremental budget constraint. After processing both days, the model assesses performance on both days with the total quantity of spares. Often the availability on the first analysis day will then exceed its target because the model typically purchases additional spares to meet the second day’s target over and above the spares required to meet the first day’s target. Whether the model is performing a one-day or two-day analysis, it produces curves of performance versus cost.

Analysis of typical scenarios, as shown in Figure 5-1, requires a judicious choice of analysis days. Often a choice of day zero as the first analysis day is

indicated, since this would give the requirement for peacetime operations. One would like to identify the "worst" or most demanding day for the second analysis day so that spares sufficient to support that day are also adequate to maintain the required availability rate throughout the 30-day wartime support period. From an activity-level perspective, day 10 might be a good choice for analysis of Aircraft A, since its flying hours drop after that, while day 30 might be appropriate for Aircraft B. But, if the scenario involves maintenance and resupply disruption until day 15, that day might be the best choice for all three types of aircraft. Trial and error is sometimes the only way to make the choice of an analysis day.

A brief description of some of the more common scenarios and ASM modes of operation follows.

5.1.2.1 SINGLE-DAY ANALYSIS

- ◆ *Steady-state conditions.* This mode typically computes spares for peacetime conditions. The single analysis day is set to zero, and a single fleet flying-hour value is specified for all days of the period.
- ◆ *Dynamic conditions.* This mode is typically used to calculate requirements for a wartime engagement only, often involving a deployment. All spares and aircraft are available at the beginning of day 1, and all pipelines are empty. The single analysis day is set to be the final day of the dynamic period, and the fleet flying-hour program for the period is specified. Steady-state flying hours (for those days preceding the period) are set to zero.
- ◆ *Dynamic conditions linked with steady state.* This option develops spares requirements for wartime conditions but assumes that steady-state conditions with peacetime item failures, pipelines, and inventory requirements are present at the start of the war. The single analysis day is at the end of the dynamic period of interest. The fleet flying-hour value for each day of the wartime period and the steady-state flying hours are specified.

5.1.2.2 TWO-DAY ANALYSIS

- ◆ *Steady-state and dynamic conditions.* This option computes requirements to support a period of steady-state operations followed by a dynamic period, (e.g., peace followed by the outbreak of war). It is implemented by a two-day analysis with the first day set to zero and the second day set to the end of the dynamic period of interest. The fleet flying-hour values for steady-state conditions and for each day of the wartime period are specified. (The ASM has the option of procuring spares for peacetime first and then procuring the additional spares to meet the wartime target next, or vice versa.)

- ◆ *Two dynamic analysis days.* This is the typical USAF mode of operation when the ASM is used to calculate requirements for deployment kits, or mobility readiness spares packages (MRSPs). The model computes the spares required for two different days in a single wartime scenario — for instance, the last day of the surge and the end of the war (for Aircraft A in Figure 5-1, this corresponds to day 10 and day 30). The model assumes that spares are required only for the wartime support period and that all spares and aircraft are available on day 1. A fleet flying-hour program for each day of the war is specified, and the steady-state flying hours are set to zero. There are two analysis days, and the model computes the requirement to meet both of those days' targets.
- ◆ *Two dynamic analysis days linked with steady state.* This mode is similar to the preceding one, except that now steady-state conditions with their own item failures and pipelines are present at the start of the war. Two analysis war days, a fleet flying-hour program for each day of the war and the steady-state flying hours value are specified. This situation is typical of a unit that fights in place at the outbreak of a war.

5.1.3 Initial Assets and User-Specified Spares Levels

The ASM can accept user-specified initial asset levels in several different ways. In the common, zero initial asset case, the ASM determines spares mixes "from scratch." At the other extreme, the ASM can evaluate the performance and cost of a specified spares mix.

Other options allow the specification of specific stockage objectives (e.g., an initial level of spares by item, already in the inventory, or a desired minimum level). These options permit the user, who has previously made stockage decisions concerning some or all items, to include those decisions in the model's solution. The user's initial levels may be items procured previously or items needed but not yet procured. The former — starting or initial assets — do not increase the costs, and the latter — user-specified buy quantities (or negotiated levels) — force the ASM to increase the spares cost. In the first case, the user may have starting assets previously paid for and may need to determine the incremental buy necessary to reach an availability target. In the second case, the user may want to force the model to buy at least a specified quantity of spares. The stock parameters allow the user to include those decisions. Furthermore, the user can specify starting assets or buys as either a minimum value, letting the model buy more if necessary, or as a fixed value, so that the model will purchase only to the user-specified level. For user-specified buy quantities, the model can handle some special cases (e.g., excluding the item from having an impact on aircraft availability, or forcing the model to buy at least a specified percentage of the item's total pipeline).¹

¹In earlier chapters, we referred to this as the mean, or expected, number of units in the pipeline. When we use the word "pipeline" here (without "mean"), it refers to the physical pipeline or the mean number of units in the physical pipeline.

5.1.4 Examples of ASM Sensitivity

In the next several sections, we present sample results to demonstrate some of the model's key capabilities. We obtained these results using items from an USAF F-15E MRSP — a “flyaway kit” of spares that deploys with the squadron. We use a modified USAF F-15E database to obtain the indenture structure, demand rates, unit costs, NRTS rates, and other item information. Parameters, scenario descriptions, and selected data elements have been changed and are illustrative only (see Table 5-1). All items undergo one of three types of depot repair — some items are repaired at a local depot, some are sent to a nearby contractor for repair, and some are ordinarily repaired at a remote source that will be cut off during war. (The third source is not typical of USAF operations, but it is common for many other countries.) For convenience, we set repair and shipping times and condemnation rates constant across items (DRT, BRT, OST, and ConPCT, respectively).

Table 5-1.
Resupply Assumptions for F-15E Spares Kit

Type of depot repair	DRT		BRT		OST (days)	ConPCT	Percentage of items
	Peace (days)	War (days)	Peace (days)	War (days)			
Type 1 – local	60	10	14	3	1	0.05	50
Type 2 – contractor	100	15	14	3	1	0.05	40
Type 3 – foreign	180	Infinite	14	3	1	0.05	10

Note: Kit contains 275 LRUs and 406 SRUs.

5.2 OPTIMIZATION WITH MULTIPLE LEVELS OF INDENTURE

Each aircraft has an indenture structure. Aircraft are composed of LRUs, LRUs are composed of SRUs, SRUs are composed of lower indenture SRUs, and so on. The ASM develops the optimal balances between procuring LRU and SRU spares. Since the indenture structure can have a great impact on the spares mix, it is important to specify that structure correctly.

Consider three different aircraft, each consisting of the same five items (with the item pipelines and costs displayed in Table 5-2), but each having a different indenture structure.

- ◆ Aircraft A — *One level* of indenture, and all five items are LRUs.
- ◆ Aircraft B — *Two levels* of indenture, with item #1 an LRU and the other four items SRUs.

- ◆ Aircraft C — *Three levels* of indenture, with item #1 an LRU, items #11 and #12 second-indenture-level SRUs, and items #21 and #22 third-indenture-level SRUs (under item #12 in the indenture structure).

Table 5-2.
Spares Mixes for Three Different Indenture Structures
(\$2,000 Investment)

Item #	Pipeline	Cost (dollars)	Aircraft A — 1 level and all LRUs (spares)	Aircraft B — 2 levels with 1 LRU (spares)	Aircraft C — 3 levels with 1 LRU (spares)
1	3.0	1,000	1	2	1
11	1.0	400	1	0	1
12	1.0	300	1	0	2
21	0.5	150	1	0	0
22	0.5	150	1	0	0

Note: In the table, "pipeline" refers to pipeline not including the AWP segment shown in Equation 4-35.

Each item's pipeline gives a rough measure of its need for stock; we explained how item backorders (and therefore aircraft availability) depend on the pipeline and the spares level in Chapters 2 and 3. Note that, for Aircraft B and Aircraft C, the sum of the costs (and the pipelines) of the lower indenture items equals the cost (and pipeline) of their parent. We have chosen the item characteristics this way so that the differences in the spares mixes for Aircraft A, B, and C reflect differing indenture structures (rather than differences in costs or pipelines).

As shown in Table 5-2, if all the items are LRUs (Aircraft A), the optimal mix for \$2,000 is one unit of each. If only item #1 is an LRU (Aircraft B), the best expenditure of \$2,000 is to buy two units of item #1. This reflects the fact that SRU backorders have a lesser effect on aircraft availability than LRU backorders. Aircraft C has the most complicated indenture structure, and the best \$2,000 spares mix reflects a compromise. SRU #12 is selected for an additional spare because it is less expensive than SRU #11 and with the addition of AWP delays for items #21 and #22 to its pipeline, has a larger pipeline; it provides a greater benefit per dollar.

As the levels of indenture increase, keeping total spares cost constant, the performance improves. Increasing the number of levels of indenture from one to two changed the number of types of LRUs in the spares mix from five to one, as shown in columns 4 and 5 of Table 5-2, and improved the ENMCS figure from 2.94 to 2.75, as shown in rows 1 and 2 of Table 5-3. In the latter case, there are fewer items that can directly ground the aircraft. Increasing the number of levels of indenture from two to three also reduces ENMCS because moving two SRUs to a third level of indenture adds another buffer layer between their failures and backorders and grounding an aircraft. However, the effect is less significant; reducing the number of LRUs, whose backorders ground aircraft,

matters more than moving some SRUs from the second to the third level of indenture.

Table 5-3.
Comparing Performance for Three Different Indenture Structures

Aircraft	Indenture structure	Total cost (dollars)	ENMCS	EBOs
A	1 level	2,000	2.94	3.18
B	2 levels	2,000	2.75	2.95
C	3 levels	2,000	2.72	2.93

In this example, the sum of the SRU costs equaled the LRU cost. If SRUs had a cost advantage, we might buy more SRUs. This is shown in Table 5-4 for two levels of indenture. In the right-hand case (columns 5 and 6), it was more economical to improve support by reducing the LRU pipeline (with more SRUs to reduce SRU delay) than to buy more LRUs. Since less expensive parts are more cost-effective in improving availability, ENMCS also improves. The opposite effect would occur if the LRU were cheaper than the sum of its SRUs.

Table 5-4.
Comparing LRU/SRU Spares Mix with Respect to Cost

Item # Pipeline		Total SRU cost equals LRU cost		Total SRU cost is less than LRU cost	
		Cost (dollars)	Spares	Cost (dollars)	Spares
1	3.0	1,000	2	1,000	1
11	1.0	400	0	200	2
12	1.0	300	0	150	2
21	0.5	150	0	75	2
22	0.5	150	0	75	2
Total cost		\$2,000		\$2,000	
Total ENMCS		2.75		2.27	

The mix of spares changes in a similar manner if we double each SRU pipeline by doubling its demand. Table 5-5 shows that the model computes a larger SRU requirement to compensate for their greater number of failures.

Table 5-5.
Comparing LRU/SRU Spares Mix with Respect to Pipeline

Item #	Cost (dollars)	Sum of SRU pipelines equals LRU pipeline		Sum of SRU pipelines is greater than LRU pipeline	
		Pipeline	Spares	Pipeline	Spares
1	1,000	3.0	2	3	1
11	400	1.0	0	2	2
12	300	1.0	0	2	2
21	150	0.5	0	1	2
22	150	0.5	0	1	2
Total cost		\$2,000		\$2,000	
Total ENMCS		2.75		3.78	

Thus spares requirements depend on the relationship between the LRU and its SRUs in terms of cost and demand. That relationship varies from LRU to LRU. Sometimes the LRU is more expensive than its SRUs combined, sometimes less expensive. Sometimes it has a greater demand, and sometimes a lesser demand than its SRUs. In our F-15E database containing only reparable items, the total SRU costs, summed across all SRUs on a particular LRU (each SRU's cost multiplied by its QPA), divided by the LRU cost, ranged from 0.07 to 10.41, averaging 0.87. In other words, while the total SRU cost averages 87 percent of the LRU cost, the total SRU cost may be much greater or may be much less than that of the LRU. Table 5-6 shows data comparing SRU cost and demand to LRU cost and demand.

Table 5-6.
F-15E Relationship Between LRU and Its SRUs

Comparison: SRUs vs. LRUs	Average	Minimum	Maximum
Sum of SRUs vs. LRU cost	0.87	0.07	10.41
Sum of SRUs vs. LRU demand	1.26	0.17	7.80
Number of SRUs on the LRU	7.80	1	25

Note: F-15E database for the 54 LRUs with SRUs.

The ASM balances spares levels based upon the various LRU/SRU relationships as well as the item characteristics. If SRUs are expensive with less demand, then fewer SRUs are selected, and vice versa. This is not an all-or-nothing situation. The best mix is not made up only of LRUs or only of SRUs but consists of some of each. Even if an SRU is inexpensive, its spares can only reduce the AWP portion of the LRU pipeline. It may be preferable to buy an

LRU that costs more than its SRUs, but which has a greater effect on aircraft availability. Thus buying a large number of inexpensive items may not be advantageous; the model must explicitly determine the proper mix between the LRUs and SRUs.

5.3 THE EFFECTS OF CANNIBALIZATION

The ability to consolidate broken LRUs onto a single aircraft greatly improves aircraft availability without increasing procurement costs. The disadvantage of cannibalization is that first-line maintenance actions increase. Cannibalizing a part requires the additional steps (and risk) of borrowing items from other NMCS aircraft and then reinstalling them later when the backorder is filled. In this section, we give examples comparing model runs with and without cannibalization and using various objective functions described in Chapter 4. We limit the discussion to the case where the model is run for a single analysis day, so that the run either does or does not use cannibalization; two-analysis-day runs are discussed in Section 5.4.

5.3.1 Cannibalization with ENMCS Optimization

In this section, we describe the effects of cannibalization under ENMCS optimization and the sensitivity of those effects to model inputs; we discuss cannibalization under other objective functions in Section 5.3.2.

We have already mentioned that cannibalization raises aircraft availability for a given cost. But it is also true that, for a given availability, as the degree of cannibalization increases, so does the number of expected backorders. To reach a given availability target, fewer spares are required, but since the number of failures is the same, there will be many failures that do not ground aircraft but still create backorders. Moving from a no-cannibalization run to one with cannibalization also reduces the bias toward buying low-cost items in the optimization process. We discuss this further in our explanation of sample model runs below.

Two model inputs with a significant effect on the number of cannibalization actions are the NMCS target and item QPAs. Larger NMCS targets permit more cannibalization; since cannibalization is free (from a procurement perspective),² the model uses it to reduce NMCS aircraft.³ To see why each item's QPA has an impact on the degree of cannibalization, consider an item with a QPA of three.

²Backorders and maintenance actions are increased by cannibalization, and may be viewed as costs.

³The use of different NMCS targets may result in different costs to achieve the same availability or ENMCS. The results of this section cannot be generalized to comparisons of all kits with regard to the projected amount of cannibalization.

This item has potentially three times as many parts available to be installed on an aircraft NMCS for the item than does an item with a QPA of one.

To demonstrate the difference between model results with and without cannibalization, we used the F-15E sample kit for a single day under steady-state conditions. We ran the model with ENMCS optimization and NMCS targets of 2, 4, and 6 (or an availability of 90, 80, and 70 percent, respectively, assuming 20 aircraft) with and without cannibalization. To reach the same target, the three cannibalization runs respectively spent 12, 39, and 64 percent less money (see “% delta \$” in Figure 5-2) than did the three runs without cannibalization. But cannibalization created 4, 221, and 372 more backorders (or more maintenance actions) for the three runs. Thus as the allowable number of NMCS aircraft increases, the model assumes that more spare parts are cannibalized from those broken aircraft, so that fewer spares are needed to reach the target; however, there is a greater penalty in the form of more maintenance actions.




Figure 5-2.
*Differences in Cost and EBOs Between Cannibalizing
 and Not Cannibalizing for Three NMCS Targets*

In the model run with a NMCS target of 4, items with the largest QPA had the largest reductions in spares levels. Each NMCS aircraft makes a QPA’s worth of spares available for cannibalization. Table 5-7 shows results from that run for the 20 items that demonstrated the largest reduction in spares levels (levels without cannibalization minus the levels with cannibalization) as well as the item’s QPA.

Table 5-7.
Reduction in Levels Resulting from Cannibalization
(20 Items with Greatest Change)

Items	1	2	3	4	5	6	7	8	9	10
Decrease in spares level	54	51	44	43	38	38	17	10	8	8
QPA	30	30	30	30	30	30	8	4	4	4
Decrease/QPA	1.8	1.7	1.5	1.4	1.3	1.3	2.1	2.5	2.0	2.0

Items	11	12	13	14	15	16	17	18	19	20
Decrease in spares level	8	7	7	7	6	6	6	6	6	5
QPA	3	4	3	4	2	2	2	2	2	2
Decrease/QPA	2.7	1.8	2.3	1.8	3.0	3.0	3.0	3.0	3.0	2.5

To approximate the number of aircraft stripped for spares, we use the decrease in the spares level divided by the QPA. This ratio estimates the number of additional aircraft *fully* cannibalized. For instance, if the item spares level changes from 10 without cannibalization to 4 with cannibalization and the item QPA is 3, then the item's ratio is 2 — a rough estimate of the additional aircraft that are stripped of this item. The ratio ranged between 1.3 and 3 for the 20 items. So, on average, cannibalization uses between 1.3 and 3 aircraft as a source of spare parts to give an ENMCS of 4.

Comparing individual component spares levels shows a tendency for optimization under cannibalization to increase the levels of high-cost spares when the total investment is kept constant. In general, of course, one reaches a given performance for a reduced investment when cannibalization is practiced, and so item levels overall tend to decrease. When investment is held constant, however, buying under cannibalization tends to reduce cost minimization's inherent bias toward low-cost items. At each buy in generating the benefit-versus-cost curve under cannibalization, the item with the largest expected backorders per QPA effectively determines availability. Buying spares of other items yields little benefit,⁴ unless the number of aircraft down for this item (the "long pole in the tent") can be reduced, whatever its cost.

Table 5-8 shows two \$40 million spares mixes, one derived allowing cannibalization, and one not. Allowing cannibalization increased the spares levels of 71 components, which were relatively high-priced (average \$115,000), while reducing the spares levels for 285 components lower cost components (average \$12,000).

⁴These items have smaller *expected* backorders per QPA than the first item, but there is still a positive probability that they have more backorders per QPA than that item.

Table 5-8.
Changes to Spares Mix Resulting from Cannibalization
(\$40 Million Investment)

Effect of cannibalization	Percentage of items	Number of items	Number of spares	Average cost of a spare (dollars)	Average QPA by item
Model buys more	10	71	86	115,000	2.0
Model buys less	42	285	846	12,000	1.5

In summary:

- ◆ Cannibalization significantly reduces the spares budget required to reach a particular availability, or improves availability for the same cost.
- ◆ As the NMCS target increases, both budget savings and the maintenance penalty of cannibalization increase relative to the no-cannibalization case.
- ◆ The mix of spares with cannibalization tends to include a few more expensive spares (with low QPA) and far fewer inexpensive spares (with high QPA) than does an *equal cost* mix derived without cannibalization.

5.3.2 Cannibalization with Confidence or ENMCS/EBO Optimization

Recall from Chapter 4 that, when cannibalization is allowed, the ASM can build spares mixes using three objective functions — confidence, ENMCS, and EBO/ENMCS:

- ◆ Confidence optimization maximizes the probability that the number of aircraft NMCS is not greater than a given target D . This probability is often called the confidence of meeting the target D , and marginal analysis to maximize this confidence is called confidence optimization.
- ◆ ENMCS “optimization” considers the confidence for all possible NMCS values and optimizes a weighted sum of them.
- ◆ EBO/ENMCS “optimization” considers two measures of performance, ENMCS and EBOs, and attempts to optimize both simultaneously.

Table 5-9 displays the results of four model runs, all using the same input and same NMCS target of 4 but each using a different objective function. In the “no cann” case, the only optimization technique is the availability optimization discussed in Chapter 2; the results of this case are displayed on the far left (that case and the ENMCS case results were described in detail in Section 5.3.1). As

one moves from left to right in the table, the degree of cannibalization increases, accompanied by a decrease in spares costs (except between the last two columns) and an increase in EBOs.

Table 5-9.
Results of Different Optimization Techniques
(NMCS Target = 4)

X-axis	No cann	EBO/ENMCS	ENMCS	Confidence
Buy cost	66.7	41.3	40.5	40.8
Confidence	64	74	74	76
EBOs	4	63	225	259

While maximizing confidence is not the same as minimizing ENMCS, the spares mix, achieved confidence, and achieved ENMCS are not strikingly different with the two optimization methods (columns 4 and 5). A spares mix built using confidence optimization normally yields a slightly higher confidence level (but also a slightly higher cost and significantly more expected backorders).

Although using the ENMCS objective produces the lowest cost solution, its maintenance workload (as indicated by the number of expected backorders) is high. Choosing the EBO/ENMCS objective (column 3) develops a compromise spares mix that costs slightly more than that resulting from using a pure ENMCS objective but yields significantly fewer expected backorders.

In Section 5.3.1, we described in detail the differences between no cannibalization and cannibalization at the item level under ENMCS optimization. The two other optimization methods (EBO/ENMCS and confidence) assume cannibalization, and produce item-level effects similar to those of ENMCS optimization. Figure 5-3 demonstrates that point. To make that comparison, we estimate, as before, the number of aircraft cannibalized by taking the item spares level assuming cannibalization subtracted from the item's level assuming no cannibalization, divided by the item QPA. We display the number of additional (compared to the "no cann" case) aircraft cannibalized under each optimization method. The bars labeled "maximum" refer to the maximum, across items, of the number of additional aircraft NMCS for an item. The bars labeled "average" refer to the average increase in the number of aircraft NMCS for an item, again taken across items.




Figure 5-3.
The Effect of Optimization Method on the Number of Aircraft Cannibalized in a 20-Aircraft Fleet

It is interesting to compare Table 5-9 with Figure 5-3. Moving from EBO/ENMCS to ENMCS to confidence as the optimization methods gives a fourfold increase in expected backorders and associated cannibalizations and maintenance, while the average number of aircraft cannibalized increases only slightly.

In the next section, we will develop a spares mix assuming a period of peace followed by a war, with no cannibalization in peace and full cannibalization in war. When a spares mix is based on two model runs (two analysis days), one of which does not allow cannibalization and one that does, the difference in results between the optimization methods is minimal. The no-cannibalization model run will buy spares with the goal of driving down EBOs, and the optimization methods used in the run with cannibalization differ primarily in their achieved EBOs; both the achieved confidence and ENMCS differ little between methods.

5.3.3 Availability Curves

It is useful to have the model display an availability-versus-cost curve showing the cost of achieving all potential availability (ENMCS) targets. Such curves are available, with some caveats. For runs without cannibalization, and for runs with cannibalization that use confidence optimization, the curves are completely optimal. That is, at each point on the curve, the cost is the minimum possible to achieve that no-cannibalization availability (or full-cannibalization confidence). The curves for an ENMCS objective and for the ENMCS/EBO objective do not guarantee optimality. Further, the spares mix depends upon the NMCS target, and the "quality" of the mix is degraded at points far from the target.

Figure 5-4 is a plot of three curves from model runs with NMCS targets of 2, 4, and 6. Notice that, in general, the curves are close and even coincide with similar ENMCS values for the same cost. However, at any cost value where the achieved ENMCS on one curve is significantly different from the target for that curve, and where the achieved ENMCS on a second curve is close to the target for that second curve, the two curves are separate. For these cost values, the curve whose achieved ENMCS is closest to its target is preferred (i.e., the second curve). Thus if one enters a NMCS target and a budget constraint, and the resulting ENMCS output is significantly different from the input NMCS target, one should rerun the model with a new NMCS target equal to the ENMCS result. By experimenting with the target, one can also cause the model to compute the spares mix that gives the lowest possible ENMCS for a specified cost.




Figure 5-4.
ENMCS-vs.-Cost Curves for Three NMCS Targets

5.3.4 Other Model Options

Several other model options affect the spares mix. One option forces the system to have at least as many spares as an item's pipeline — a simple way of reducing backorders when assuming cannibalization. Another option forces the model to buy at least as many spares as the pipeline for high-QPA items, which under cannibalization normally would not require spares. All options have their

advantages and disadvantages, and all of them can be used to tailor the model to the needs and policies of the logistics system being modeled.

5.4 DYNAMIC FLYING-HOUR PROGRAMS

To determine requirements or evaluate support for a dynamic scenario such as war, the analysis should focus on the logistically most demanding day. Usually that is the last day of the wartime support period; however, for scenarios with an initial surge or a significant drop in flying hours over time, the best analysis day is less certain. To demonstrate that point, we ran the model under a typical scenario with our demonstration database. The flying-hour profile started with steady-state conditions of 10 hours a day, moved to a surge period of 60 flying hours a day, and then leveled off at 40 hours a day after day 5. Figure 5-5 displays the flying-hour profile and the resulting ENMCS plotted for various days of the war. ENMCS worsened over the surge period and then improved (dropped) once the surge ended. With this scenario, the lowest availability was projected to occur at the end of the surge period (day 5), making that day a good choice for the analysis day; even with flying hours dropping by a third, the ENMCS results varied only by half an aircraft or a few percent in availability after the fifth day.




Figure 5-5.
ENMCS Over a Period of Dynamic Flying Hours

5.5 SPARES MIXES FOR PEACE AND WAR

Sometimes a user wants to develop a spares mix for peace and war, as described in Section 5.1.2. The flying-hour profile is originally steady state, followed by a dynamic period during which flying hours vary day by day.

To develop a peacetime and wartime spares mix, the user must enter the availability or budget targets for each analysis period and then run the model. The model purchases enough spares to meet the first target and next purchases additional spares to meet the second target. After processing both days' parameters, the ASM returns to the first analysis day and recalculates performance with the total quantity (a shopping list) of spares purchased to support both periods.

In other words, the ASM does not perform a simultaneous integrated spares computation but instead does one day at a time. Thus there is some flexibility in operating the model. One can first develop a mix of spares for peace and then add to that mix to support a wartime scenario, or one can develop a spares mix for war first (still assuming the peacetime flying hours) and then add to that mix to reach the peacetime availability target. Although this order is counterintuitive (since it is not chronological), it may produce a better spares mix. The recommended way to operate the model (peace, then war — or war, then peace) is to run the dominating or most demanding day first (i.e., the day on which it costs the most to achieve a target availability).

Again consider the F-15E sample kit (see Table 5-1). Suppose that the 20-aircraft fleet flies 10 hours per day and that maintenance does not perform cannibalization in peacetime. For wartime operations, assume 40 flying hours on day 1, 60 hours per day for the rest of the war, and that maintenance does cannibalize.

Figure 5-6 displays the availability-versus-cost curve for peace and war, with peacetime availability for day 0 and wartime availability for day 24. If the total budget is spent preparing for peacetime operations, the peacetime curve in Figure 5-6 is the best result the model can obtain. The wartime curve also represents the best results the model can obtain, although as explained earlier, the wartime curve is really composed of several curves produced by several model runs assuming different NMCS targets, and is only approximately optimal. The same dollars produce much higher wartime than peacetime availability (i.e., peace is the more demanding period) for expenditures under \$70 million, mostly because the wartime scenario assumes cannibalization, so there is an additional source of spares in war. After the \$70 million point on the curve, peacetime performance exceeds that of war because, for high availabilities, little cannibalization can be done and other factors become the drivers (e.g., greater flying hours or repair times).

5.5.1 Meeting Availability Targets for the Lowest Cost

In developing a peacetime and wartime spares mix, the relationship of the individual peacetime and wartime curves is critical. At an availability that costs more to achieve in peace than war (peace dominates war), the spares mix for peacetime should be laid in first to focus on reaching the more difficult target for the least cost. Suppose the target is 80 percent availability for both peace and war. Peace dominates war, since peacetime operations require \$66.7 million to reach the 80 percent availability target, while wartime operations require only \$51 million (as shown by Figure 5-6).




Figure 5-6.
Peacetime and Wartime Availability — Cost Curves Plotted Separately

We then take the \$66.7 million in assets needed to reach 80 percent availability (see Table 5-10) in peacetime operations and see that those assets do support wartime operations. We evaluate the \$51 million in assets, which is enough to maintain 80 percent availability in wartime operations, and find that it is insufficient to achieve 80 percent availability in peacetime. In fact, to do so, the mix must be supplemented by \$15.7 million. Thus the first option produces a lower cost mix (by \$1.5 million) and supports both wartime and peacetime operations at the required availability.

Table 5-10.

Example: When Peace Dominates, with an 80 Percent Availability Input Target for Peace and War

Type of model run	Achieved availability (percent)		Cost (\$ millions)		
	Peace	War	Peace	War	Total
Peace, then war	80	80	66.7	0	66.7
War, then peace	80	85	17.0	51	68.2

5.5.2 Meeting a Fixed Budget Constraint

Suppose the same scenario as before, but now with a budget constraint of \$50 million. The question is how to split that \$50 million between peace and war. As with the previous example, we first bound the problem by determining the consequences of allocating the entire \$50 million to develop the optimal mix for peace or, alternatively, allocating it all to best support wartime operations. Table 5-11 displays those extremes, which greatly reduce the options. For this example, unless the wartime scenario is solely important, it would be unwise to put all of the money there, because peacetime availability drops to zero. On the other hand, the peacetime spares mix performs well in war. By first bounding the problem and then running a few compromise solutions that vary the budget split between war and peace, the user should be able to achieve a suitable compromise spares mix.

Table 5-11.

Example: Bounding the Peace and War Availabilities for a \$50 Million Budget

Type of model run	Achieved availability (percent)		Cost (\$ millions)		
	Peace	War	Peace	War	Total
Peace, then war	45	69	50	0	50
War, then peace	0	80	0	50	50

6.0 Initial Provisioning

A version of the ASM determines spares requirements in a special case — initial provisioning. Much of the development was done under the sponsorship of the Israel Air Force, and resulted in a variant of the ASM called ISAAC — Initial Spares Aircraft Availability Calculation. Some of the adaptations we made to the model to handle initial provisioning are discussed in this chapter; other matters concerning the actual model operations are documented in a forthcoming LMI report [Kline, et al.]. Initial provisioning is the process of procuring spares in connection with the procurement of a new weapon system, or with procurement of additional units of a weapon system already in the inventory. The purpose of these spares is to support the new system over some initial “coverage” period until the standard spares replenishment system can begin to provide support. The ASM addresses three basic questions pertaining to initial provisioning: Which spares should be ordered within a specific coverage period (given a constrained budget or availability target)? When should they be ordered? How should funds from the total budget be obligated over time?

The major changes required to address the coverage period are in accounting for condemnations and estimating annual budgets for the coverage period. We must also incorporate common components (i.e., the economies of scale made possible when the aircraft undergoing initial provisioning has parts in common with previously procured aircraft).

6.1 INITIAL PROVISIONING TIME FRAME

Figure 6-1 presents an example of the model’s approach for estimating initial spares procurement. The figure displays two key periods:

- ◆ *Coverage period* — the period from the time the first aircraft arrives to the time the standard replenishment system can begin to provide spares support.
- ◆ *Ordering and budget period* — the period in which to allocate funds and order the spares so they arrive on time. (We assume that the user incurs costs and must provide budget authority when the order is placed.) An annual procurement equals the cost of all orders placed in a given year (the buy cost).




Figure 6-1.
Time Frame for Initial Provisioning

Figure 6-1 shows an initial provisioning coverage period of 3.5 years with the aircraft delivered in 1998. Thus the coverage period begins on the first day of the year the aircraft are to be delivered (1998) and ends 3.5 years later, or midway through 2001. The model estimates the aircraft availability at the end of the coverage period, so we must procure enough spares to provide support over that period.

The user specifies the coverage period by entering its length and the year that the aircraft will be delivered. The model determines the initial provisioning ordering and budget period using the aircraft delivery year and each item's procurement lead-time total (PLTT). We initially assume that all spares arrive in the aircraft delivery year and that their impact on budgets is one PLTT in advance of their delivery, at the time the order is placed (as is shown in Figure 6-1). The start of the ordering and budget period occurs a maximum component PLTT before delivery of the aircraft, and it ends the year before aircraft delivery. In our example, if the longest PLTT of all the items is three years, the ordering and budget period begins in 1995 and ends in 1997.

6.1.1 Flying Hours in a Coverage Period

Figure 6-2 is a nominal example of a flying-hour scenario. In this example, the only year in which flying hours change is the delivery year (the hours change over that year because not all of the aircraft are delivered simultaneously). We make a key simplifying assumption that aircraft delivery is spread out evenly over the year. Thus, the first year has only half as many cumulative flying hours as each of the following years. The result is that the number of years of flying is the coverage period less half a year (what we will term a "steady-state coverage period"). The user may also specify a war (dynamic period) occurring after the

coverage period. In that case, the model will compute the spares requirements for the war and add those to the requirements generated to support the aircraft in the coverage period.




Figure 6-2.
Nominal Flying-Hour Program for Initial Provisioning

6.1.2 Condemnations in a Coverage Period

In Chapter 3's description of pipelines for steady state and dynamic flying hours, the model looks back a PLTT to calculate the condemnation pipeline. Such is the case for three items with PLTTs of 2, 3, and 4 years in the top of Figure 6-3. Since PLTT is not usually equal to the coverage period, this pipeline must be calculated carefully. For items with PLTT greater than the coverage period, the number of items in the condemnation pipeline is only those corresponding to condemnations in the coverage period.¹ For items whose PLTT is less than or equal to the coverage period, the mean pipeline is computed, as usual, as the mean number of condemnations in the PLTT. If PLTT is strictly less, there may be condemnations early in the coverage period not included in the pipeline. These early condemnations must be included as a requirement.

¹The ASM actually computes this by setting PLTT equal to the coverage period (for this computation only).




Figure 6-3.
How Coverage Period Adjusts PLTT

We procure spares to cover these condemnations as a sunk cost without applying any marginal analysis criterion or calculating any safety level for it. (Without this buy, the projected spares level would be negative.) This fixed buy, together with the projected condemnations in the PLTT, thus accounts for all the projected condemnations in the coverage period. Safety level is calculated only for condemnations in the PLTT.

An "operating requirement" is based on the desired availability at the end of the coverage period. This, in turn, is based on the pipelines at that point (base repair, order and ship, depot repair, and condemnation) and the associated safety level. The early condemnation requirement is added to the operating requirement to obtain the total provisioning requirement. Note that there is no need for a safety level requirement for the early condemnations. If they vary from the expected number, we will have that information a PLTT before the end of the coverage period and can react by adjusting the procurement, if necessary.

6.2 COMMON COMPONENTS

Common components are components common to both the initial provisioning aircraft and other aircraft already procured. For example, a component on the initial provisioning aircraft, series F-15X, might also be on the existing F-15C and F-16D. Treatment of common components must apply any surplus stock already in the inventory toward the requirement and must consider economies of scale.

Since stock is already available for other aircraft, the new aircraft will need less inventory than they would if the component were not common. The ASM uses a simple approximation to incorporate those benefits by calculating the number of "free assets": assets available to the weapon system without procurement cost because of its commonality.

Table 6-1.
Treatment of Common Component

Basis of spares levels	Pipeline	Safety stock (calculated)	Spares requirement
F-15C and F-16D (existing)	16	—	28 ^a
New F-15X	9	4.5	14
Without considering common components	—	—	42
Considering common components	25	7.5	33

Note: Safety stock is 1.5 standard deviations, or 1.5 times the square root of the mean (pipeline).

^aAssets for F-15C and F-16D consist of pipeline quantity of 16 plus 12 additional spares.

Suppose that a new item being procured during initial provisioning for the F-15X has a pipeline of 9. That item already exists in the inventory for the F-15C and F-16D, with a pipeline of 16 and an asset position of 28. We assume that the variance-to-mean ratio of demand is 1, and that the model usually requires about 1.5 standard deviations of the pipeline quantity for a safety level, so the model spares requirement for the new item alone, without considering commonality, is 14 spares $\left(9 + 1.5\sqrt{9} \text{ rounded} \right)$. Alternatively, the total model spares requirement for the common component with a comparable safety level is the sum of the two pipelines plus 1.5 standard deviations, or $33 \left[9 + 16 + \left(1.5\sqrt{9 + 16} \right)\right]$. In contrast, if the model did not consider common components, the total inventory requirement would be 42 (28 existing assets and 14 initial provisioning items) instead of 33. Thus, the common component's economy of scale generates 9 free spares $(42 - 33)$, reducing the initial provisioning procurement to only 5 spares and bringing the total to 33 (28 assets and the cost of 5 spares). Therefore, we give the F-15X common component 9 free spares, correctly approximating the impact of common components, and permitting them to be considered properly in budget and availability calculations.

Bibliography

- Crawford, Gordon. *Palm's Theorem for Nonstationary Processes*. Santa Monica, Cal.: The RAND Corporation, R-2750-RR, October 1982.
- Feller, William. *An Introduction to Probability Theory and Its Applications*, Vol. 1., 3rd Edition. New York: John Wiley and Sons, Inc., 1968.
- Feller, William. *An Introduction to Probability Theory and Its Applications*, Vol. 2, 2nd Edition. New York: John Wiley and Sons, Inc., 1971.
- Gaver, Donald P., Karen E. Isaacson, and John B. Abell. *Estimating Aircraft Recoverable Spares Requirements with Cannibalization of Designated Items*. Santa Monica, Cal.: The RAND Corporation, R-4213-AF, 1993.
- Hillestad, Richard J., and Manuel J. Carillo. *Models and Techniques for Recoverable Item Stockage When Demand and the Repair Process Are Non-Stationary — Part I: Performance Measures*. Santa Monica, Cal.: The RAND Corporation, N-1482-AF, 1980.
- Hillestad, Richard J. *Dyna-METRIC: Dynamic Multi-Echelon Technique for Recoverable Item Control*. Santa Monica, Cal.: The RAND Corporation, R-2785-AF, 1982.
- Isaacson, Karen E., et al. *Dyna-METRIC Version 4: Modeling Worldwide Logistics Support of Aircraft Components*. Santa Monica, Cal.: The RAND Corporation, R-3389-AF, 1988.
- King, Randall M. LMI Working Note AF401-3. *Assessing Aircraft Spares Support in a Dynamic Environment*. July 1985.
- Kline, Robert C., et al. LMI Report AF502RD1. *Initial Spares Aircraft Availability Calculation — ISAAC: User's Manual*. (to be released)
- Kotkin, Meyer H. *Operating Policies for Non-Stationary Two-Echelon Inventory Systems for Repairable Items*. Aberdeen Proving Ground, Md.: U.S. Army Materiel Systems Analysis Activity, AMSAA Special Publication No. 39, 1986.
- Kruse, Karl. *An Exact N Echelon Inventory Model: The Simple Simon Method*. Philadelphia, Pa.: U.S. Army Inventory Research Office, ADA067762, 1979.
- Miller, Louis W., and John B. Abell. *DRIVE (Distribution and Repair in Variable Environments)*. Santa Monica, Cal.: The RAND Corporation, R-4158-AF, 1982.

- Muckstadt, J. A. "A Model for a Multi-Item Multi-Echelon Inventory System", *Management Science*. Vol. 20, No. 4, December 1973, pp. 472 – 481.
- O'Malley, T. J. LMI Report AF201. *The Aircraft Availability Model: Conceptual Framework and Mathematics*. June 1983.
- Papoulis, Athanasios. *Probability, Random Variables, and Stochastic Processes*, 2nd Edition. New York: McGraw Hill, 1984.
- Parzen, Emanuel. *Stochastic Processes*. New York: Holden-Day, 1962.
- Ross, Sheldon. *Applied Probability Models with Optimization*. San Francisco, Cal.: Holden-Day, 1970.
- Sherbrooke, Craig C. "METRIC: "A Multi-Echelon Technique For Recoverable Item Control", *Operations Research* 16. 1968, pp. 122 – 141.
- Sherbrooke, Craig C. "An Evaluator for the Number of Operationally Ready Aircraft in a Multi-Level Supply System", *Operations Research* 19. 1971, pp. 618 – 635.
- Sherbrooke, Craig C. LMI Report AF601LN1. *Using Sorties vs. Flying Hours to Predict Aircraft Spares Demand*. (to be released.)
- Simon, Richard M. "Stationary Properties of a Two Echelon Inventory Model for Low Demand Items", *Operations Research* 19. 1971, pp. 761 – 777.
- Slay, F. Michael. LMI Working Note AF301-3. *Vari-METRIC: An Approach to Modeling Multi-Echelon Resupply When the Demand Process Is Poisson with a Gamma Prior*. March 1984.

APPENDIX A

Marginal Analysis

Marginal Analysis

Marginal analysis is applied in several ways in the Aircraft Sustainability Model (ASM) to develop optimal spares mixes. In this appendix, we present a general proof of the optimality of marginal analysis and demonstrate its application to several problems in developing optimal inventory mixes.

Let $[n(i)]_{i=1}^m$ denote a particular array of nonnegative integers, which we will call the initial level. In relation to the initial level, define the set $\{(s_i)\}$ of all integer arrays (s_i) with the property that $s_i \geq n(i)$ for each i . We will call each such array a level. Let $\{c_i\}$ for $i = 1, \dots, m$ denote a set of positive real numbers, which we will call *costs*. For any given level (s_i) , we can define its *total cost* C , in relation to the initial level, by the equation

$$C = \sum_i [s_i - n(i)] \times c_i.$$

Next, suppose that we have functions f_i such that $f_i(s_i)$ is defined for all possible values of s_i and such that the difference functions

$$\partial_i(s_i) = f_i(s_i) - f_i(s_i - 1)$$

are all positive and decreasing:

$$0 \leq \partial_i(n + 1) \leq \partial_i(n)$$

for all n .

For all i and all n , define *sort values* $\{v_i(n)\}$ by

$$v_i(n) = \frac{\partial_i(n)}{c_i}.$$

Form the ordered list L consisting of the $v_i(n)$ in descending order. Let L_C denote any initial section of the list L , where C is the sum of the costs c_i that appear in the section. Define the level (s_i) by $s_i = m_i$, where m_i is the maximum value of n appearing in the sort values $v_i(n)$ in the sublist L_C .

We claim that the level (s_i) , as defined above, has total cost C and that, if (s_i') denotes any other level with total cost equal to or less than C , then

$$\sum_i f_i(s_i') \leq \sum_i f_i(s_i). \quad [\text{Eq. A-1}]$$

In other words, for the various possible total costs defined by the initial sections of the list L , the levels (s_i) represent undominated solutions to the problem of maximizing the sum

$$\sum_i f_i(s_i)$$

for cost C .

PROOF

Decreasing differences and the ordering of L (and therefore L_C) ensure that $v_i(j)$ is included in L_C for each j such that

$$n(i) + 1 \leq j \leq m_i.$$

This means that for each i there are exactly $m_i - n(i)$ elements in the list L_C and therefore

$$\sum_i [s_i - n(i)] c_i = \sum_i [m_i - n(i)] c_i = C.$$

So the level (s_i) does have total cost C .

Now let (s_i') be any other level with total cost

$$\sum_i [s_i' - n(i)] c_i \leq C.$$

Let A_C denote the set of sort values associated with the level (s_i') defined as follows:

$$A_C = \{v_i(k) : n(i) + 1 \leq k \leq s_i', i = 1, \dots, m\}.$$

From the construction of L_C and A_C it follows that:

$$\sum_i \{f_i(s_i) - f_i[n(i)]\} = \sum_{L_C} v_i(j) c_i, \quad [\text{Eq. A-2}]$$

and

$$\sum_i \{f_i(s_i') - f_i[n(i)]\} = \sum_{A_C} v_i(k) c_i. \quad [\text{Eq. A-3}]$$

We can rewrite

$$\sum_{A_C} v_i(k) c_i$$

as

$$\sum_{A_C} v_i(k) c_i = \sum_{A'_C} v_i(k) c_i + \sum_{A''_C} v_i(k) c_i,$$

where A''_C is the set of sort values common to A_C and L_C , while A'_C is the set of sort values in A_C but not in L_C . Since L_C is an initial section of the list L , it follows that

$$\max_{A'_C} v_i(k) \leq \min_{L_C} v_i(j).$$

Thus

$$\sum_{A_C} v_i(k) c_i \leq \min_{L_C} v_i(j) \sum_{A'_C} c_i + \sum_{A''_C} v_i(k) c_i. \quad [\text{Eq. A-4}]$$

Since the total cost of (s_i') is less than or equal to C , it follows that

$$\sum_{A'_C} c_i \leq \sum_{L'_C} c_i, \quad [\text{Eq. A-5}]$$

where L'_C is the set of sort values that are in L_C but not in A_C . Thus, Equation A-4 becomes

$$\begin{aligned} \sum_{A_C} v_i(k) c_i &\leq \min_{L_C} v_i(j) \sum_{L'_C} c_i + \sum_{A''_C} v_i(k) c_i \\ &= \sum_{L'_C} \left[\min_{L_C} v_i(j) \right] c_i + \sum_{A''_C} v_i(k) c_i \\ &\leq \sum_{L'_C} v_i(j) c_i + \sum_{A''_C} v_i(k) c_i \\ &= \sum_{L_C} v_i(j) c_i. \end{aligned}$$

Applying this last inequality to Equations A-2 and A-3 yields Equation A-1:

$$\sum_i f_i(s_i') \leq \sum_i f_i(s_i),$$

completing the proof.

Note that not only are marginal analysis solutions optimal for obtaining maximum returns for a given cost, they are also optimal for identifying minimum cost to achieve a given return. To see this, note that for any level (s_i') with total cost strictly less than C , the inequality in Equation A-5 would be strict, making the inequality in Equation A-1 strict, which means that the cost C is minimal.

Note also that marginal analysis can be applied to minimizing a sum subject to a cost constraint. In that case, the difference functions $\partial_i(s_i)$ must be negative, and the differences must be increasing

$$\partial_i(n+1) \geq \partial_i(n).$$

The problem becomes the following. For a given total cost C , defined by some initial section on the ordered list of sort values, find the level (s_i) with the property that

$$\sum_i f_i(s_i) \leq \sum_i f_i(s_i'),$$

where s_i' denotes any other level with total cost $\leq C$. By taking sort values defined by

$$v_i(n) = -\frac{\partial_i(n)}{c_i},$$

the argument goes through as before, but with the appropriate sign changes.

Note that marginal analysis produces solutions that are optimal, but may not produce all possible optimal solutions. The total costs defined by the initial sections on the ordered lists L represent a discrete set of possible values, and for those values the marginal analysis method yields optimal solutions. For intervening cost values, however, marginal analysis does not produce solutions. Finding optimal solutions for *any* specified cost would represent a version of the well-known knapsack problem. Practical applications to inventory problems suffer very little from this theoretical drawback. The set of solutions defined by the initial sections of the sort value lists is sufficiently rich to cover the full range of possible costs, expected backorder levels (EBO), and aircraft availability rates.

Application 1 — Minimizing Expected Backorders at a Single Base

Suppose we have a single location and a given set of components, $i = 1, 2, 3, \dots, I$, each with procurement cost c_i . We suppose that pipelines are determined by item characteristics and the operating scenario. We may consider the steady-state situation or the situation on a given day t in a dynamic scenario. We suppose also that depot spares levels are fixed, so our problem is to determine least-cost base-level mixes (s_i) to minimize $\sum EBO(i, s_i)$, the sum of expected backorders for component i with spares level s_i , subject to the total cost $C = \sum c_i s_i$. In the notation of the previous section, let $f_i(s) = EBO(i, s)$. Typically, but not necessarily, we start from a zero base [i.e., the initial assets $n(i) = 0$ for each i].

Let $p(i, x)$ be the probability distribution of the total base resupply pipeline for component i .

Then

$$EBO(i, s) = \sum_{x=s+1}^{\infty} (x - s)p(i, x).$$

The EBO reduction from adding the s th spare of component i to the inventory is

$$\begin{aligned} EBO(i, s - 1) - EBO(i, s) &= \sum_{x=s}^{\infty} [x - (s - 1)]p(i, x) - \sum_{x=s+1}^{\infty} (x - s)p(i, x) \\ &= [s - (s - 1)]p(i, s) + \sum_{x=s+1}^{\infty} [x - (s - 1) - (x - s)]p(i, x) \\ &= p(i, s) + \sum_{x=s+1}^{\infty} p(i, x) \\ &= \sum_{x=s}^{\infty} p(i, x). \end{aligned}$$

Thus expected backorder reductions are decreasing — i.e., the differences are negative and decreasing in magnitude. Therefore, we can apply the minimizing form of marginal analysis to show that forming a shopping list ranked in descending order of expected backorder reduction divided by cost yields initial segments that are undominated. That is, the spares levels corresponding to any such initial segment provides the lowest expected backorder total for its cost, and conversely.

Application 2 — Minimizing Total Base Expected Backorders Over All Bases in a Multi-Echelon System

This application is similar to the first one, except that our decision now involves distributing spares between base and depot as well, and we are interested in the total of systemwide expected backorders over all bases, assuming that demand is equally distributed across bases. In this case, we let $EBO(i, n)$ be the total — over all bases — of expected backorders for component i corresponding to a total inventory of n spares of component i , *distributed optimally between bases and depot*. The determination of this distribution, and thus the value of $EBO(i, n)$, is a subproblem that must be solved by comparing alternative distributions, using the approach outlined in the text. In this case, we cannot apply marginal analysis directly to $\Sigma EBO(i, n)$, since the functions $EBO(i, n)$, may not be convex (i.e., have differences of decreasing magnitude). Often the optimal base/depot distribution evolves by adding additional depot spares with minimal EBO reduction until a threshold is reached. Then, an additional spare allows a “flush out” of many spares to base level and a corresponding dramatic reduction in EBO. This phenomenon makes the final spare — or sometimes the next spare — seem more valuable than those immediately previous. This

situation is handled by treating such spares as a group and pooling their benefits and costs. Geometrically, we replace the function $EBO(i, n)$ with its convex hull, eliminating the "spurious" intermediate points. We can then apply marginal analysis as before.

Application 3 — Maximizing Aircraft Availability (Without Cannibalization)

Suppose we wish to maximize aircraft availability in a multi-echelon environment (either in a steady state environment or at a specific day t in a dynamic scenario) for a fleet of aircraft composed of components $i = 1, 2, 3, \dots, I$, all of which are line replaceable units (LRUs). Suppose also that we wish our spares mix not to rely on cannibalization. Let the expected backorder functions $EBO(i, n)$ represent the total expected backorders for component i with n spares distributed optimally between bases and the depot. Let $A_i(n)$ represent the probability that a random aircraft is *not* missing a unit of component i , as in Equation 2-20. Then we wish to find spares levels (s_i) that maximize the availability

$$A = \prod_i A_i(s_i),$$

subject to a cost constraint on $C = \sum_i c_i s_i$.

To apply the maximizing form of marginal analysis, we consider $\ln A = \sum \ln[A_i(s_i)]$. This is a separable sum of increasing functions, and maximizing $\ln A$ will, of course, maximize A . The differences $\delta_i(s_i) = \ln A_i(s_i) - \ln A_i(s_i - 1)$ are positive. However, as in Application 2, the multi-echelon tradeoff between base and depot may result in $EBO(i, n)$ being nonconvex, which, in turn, results in $\ln[A_i(n)]$ being nonconcave. That is, the differences may not be monotonically decreasing. We can apply the approach of the previous section (i.e., take the concave hull of each $\ln A_i(n)$ by grouping spares together when nonconcavities occur to form a super-spare with an aggregate cost and an average benefit), providing a modified objective function to which we can apply marginal analysis. Thus, forming the shopping list L in order of

$$V_i(n) = \frac{\ln A_i(n) - \ln A_i(n_i - 1)}{C_i}$$

provides undominated solutions. [This is true with the restriction that we force component sort values to be decreasing by grouping them to form the concave hull of $A_i(n)$. In practice, the only penalty one pays for this is a slight increase in the "graininess" of the availability-versus-cost curve — some of the cost and availability increments between entries have increased.]

To incorporate shop replaceable units (SRUs) into this solution, we take the product above over the LRUs only. However, we first solve a subproblem within each LRU and its family of SRUs, so that $EBO(i, n)$ now represents the

minimal expected backorders for LRU i with an expenditure of $n \times C_i$ divided optimally between spares of LRU i (with optimal base-depot distribution) and its SRUs. (See Sections 3.6.2 and 4.3.)

APPLICATION 4 — Optimizing Confidence Level

Suppose, at a single location, we wish to find a least-cost spares mix that maximizes the probability that fewer than $T > 0$ aircraft are down [that is, not mission-capable-supply (NMCS)] for lack of spares, and we allow cannibalization as an intrinsic element of our sparing strategy. Typically, we are interested in this problem at a particular point in time during a dynamic wartime scenario, but it may be addressed similarly in a steady-state context. For simplicity, we will assume that the quantity per application is 1 and that the application percentage is 1.0 for all LRUs i on the aircraft.

As in Chapter 4, we let X_i be the random variable of the number of components i in the resupply pipeline. Then, as in Equation 4-5, the probability of fewer than T aircraft being NMCS is given by

$$\begin{aligned} p(\text{NMCS} \leq T) &= \prod_i p(\text{BO}_i \leq T) \\ &= \prod_i p(X_i \leq s_i + T), \end{aligned}$$

where BO_i represents the backorders for component i , s_i represents the spares level for component i , and the product is taken over all LRUs on the aircraft.

As in Application 3, we can take the logarithm to express the objective function as a sum of independent functions,

$$\begin{aligned} \ln[p(\text{NMCS} \leq T)] &= \sum_i \ln p(\text{BO}_i \leq T) \\ &= \sum_i \ln p(X_i \leq s_i + T) \end{aligned}$$

and proceed as before.

APPENDIX B

Computing Backorder Statistics: An Example

Computing Backorder Statistics: An Example

ASSUMPTIONS

Our measure of overall supply performance is aircraft availability. In the dynamic case, we must also specify on which day we measure availability. It is frequently the case that availability must be examined on several days to capture a true picture of the dynamic situation. To estimate availability, we need to approximate the probability distribution of base backorders on those days. As we saw in Chapter 3, the probability distribution for backorders on a given day is easily obtained from the probability distribution for the number of items in the base resupply pipeline on that day.

This example illustrates how the model calculates pipeline and backorder probability distributions for day 6 of a dynamic scenario. We begin with no stock at the depot, and then show how the base backorder probability distribution function (PDF) changes as we add spares at the depot.

We assume that the resupply times and failure rate per flying hour are constant (but different) for steady-state and dynamic conditions. All model computations and outputs are defined as of the end of each day's flying activity. We assume the flying-hour program (FHP) shown in Table B-1, and we assume that demands arise from a nonhomogeneous Poisson process driven by flying hours, as in Equation 3-26. Table B-1 describes a scenario with a five-day surge at six times the peacetime rate, followed by a sustained period at four times the peacetime rate. The first 15 days of this scenario are shown in Figure 3-1.

Table B-1.
Sample Scenario

Day	0	1	2	3	4	5	6
Flying hours per day	100	600	600	600	600	600	400

Consider a particular line replaceable unit (LRU) with no subassemblies and item characteristics shown in Table B-2. Half of all failures are repaired at the depot level, the others are base repairable, there are no condemnations, and the failure rate is 0.01 per flying hour. We also assume that both the quantity per application (QPA) and the future application percentage (FAP) for this item are 1.

Table B-2.
Item Information with Constant Resupply

Item characteristics	Item information	
	Peacetime [steady state ($t \leq 0$)]	Wartime [dynamic ($t > 0$)]
Base repair time (BRT)	5 days	5 days
Order and ship time (OST)	3 days	3 days
Depot repair time (DRT)	10 days	10 days
Failure factor = $FF \times FAP \times QPA$	0.01 per flying hour	0.01 per flying hour
Not reparable this station = NRTS	0.5%	0.5%
Condemnation rate = ConPCT	0%	0%

PIPELINE STATISTICS

We begin by using Equation 3-27 to compute the mean of the base repair pipeline on day zero:

$$\begin{aligned}
 BRpipe(T) &= \sum_{k=T-BRT+1}^T \lambda(k)[1 - NRTS(k)] \\
 &= \sum_{k=0-5+1}^0 [FF(k) \times FHP(k) \times QPA \times FAP][1 - NRTS(k)] \\
 &= \sum_{k=4}^0 (0.01 \times 100 \times 1 \times 1)(1 - 0.5) = 5 \times (0.01 \times 100) \times 0.5 = 2.5.
 \end{aligned}$$

Because of our assumption that demand is a nonhomogeneous Poisson process, we have $VBRpipe(0) = BRpipe(0) = 2.5$.

On day 6, the computation is similar, except that the flying hours vary by day:

$$\begin{aligned}
 BRpipe(6) &= \sum_{k=6-5+1}^6 [0.01 \times FHP(k)](1 - 0.5) \\
 &= 6 \times 0.5 + 6 \times 0.5 + 6 \times 0.5 + 6 \times 0.5 + 4 \times 0.5 \\
 &= 14; \\
 VBRpipe(6) &= BRpipe(6) \\
 &= 14.
 \end{aligned}$$

Simply put, to obtain the mean of the base repair pipeline on day T , accumulate the mean base repair inductions (mean failures multiplied by the base repair rate) over the BRT days that end on day T . Note that a period of BRT days *ending* on day T encompasses days $T - BRT + 1$ through T . $BRpipe(T)$ as a function of T is as shown in Figure B-1.




Figure B-1.
Mean of Base Repair Pipeline for the Example

Calculating the mean (and variance) of the *OSpipe* is analogous to calculating the base repair pipeline: we accumulate the mean depot repair inductions over the interval $(T - OST, T)$. From Equation 3-28 and the fact that $OST = 3$, we have

$$OSpipe(6) = \sum_{k=6-3+1}^6 FF(k) \times FHP(k) \times NRTS(k) = 6 \times 0.5 + 6 \times 0.5 + 4 \times 0.5 = 8.0,$$

and because of our nonhomogeneous Poisson demand assumption, we have $VOSpipe(6) = OSpipe(6) = 8.0$.

Next we evaluate, for day T , expected depot backorders and the variance of depot backorders at time $T - OST$. The reason for looking back an OST is that conditions at the depot affect the base an OST later. (In our example, the depot could disappear on day 4 and the base would not be affected on day 6!)

First we use Equation 3-29 to compute the mean and variance of the depot repair pipeline. Letting $T = 6$ and recalling that $OST = 3$, we have

$$DRpipe(T - OST) = DRpipe(3) = \sum_{k=6-3-10+1}^{6-3} FF(k) \times FHP(k) \times NRTS(k) = 12.5.$$

BACKORDER STATISTICS WITH NO DEPOT STOCK

Recall that the number of depot backorders in which we are interested is the number at time $T - OST$. Since we assume for the moment that the depot stock level $s_0 = 0$, the number of depot backorders is equal to the number in the depot repair pipeline. Therefore $DEBO(T - OST) = DEBO(3) = DRpipe(3) = 12.5$. Because of our assumption that the demand process is nonhomogeneous Poisson, the depot repair pipeline at time $T - OST$ is Poisson, and we have

$$VDBO(T - OST) = VDBO(3) = VDRpipe(3) = DRpipe(3) = 12.5.$$

Now we use Equations 3-30 and 3-31 to compute the mean and variance of the base resupply pipeline at time $T = 6$:

$$Bpipe(6) = BRpipe(6) + OSpice(6) + DEBO(3) = 14 + 8 + 12.5 = 34.5,$$

$$VBpipe(6) = VBRpipe(6) + VOSpipe(6) + VDBO(3) = 14 + 8 + 12.5 = 34.5.$$

(In general, the mean and variance of the base resupply time would be distinct.)

From Equation 3-14, we have $VMR(6) = \frac{\sigma^2(6)}{\mu(6)} = \frac{34.5}{34.5} = 1$, so that the probability distribution for the base resupply pipeline when the depot stock $s_0 = 0$ is the Poisson distribution rather than the negative-binomial distribution in Equation 3-15,

$$\Pr(BpipeRV_6 = k) = \frac{e^{-\mu(6)} \mu(6)^k}{k!} = \frac{e^{-34.5} (34.5)^k}{k!}.$$

If the base stock level is, say, $s = 30$, then the probability of no base backorders on day 6 is

$$\Pr(BORV_6 = 0) = \Pr(BpipeRV_6 \leq 30) = \sum_{k=0}^{30} \frac{e^{-34.5} (34.5)^k}{k!} = 0.253.$$

The probability of exactly 5 backorders is

$$\Pr(BORV_6 = 5) = \Pr(BpipeRV_6 = 35) = \frac{e^{-34.5} (34.5)^{35}}{35!} = 0.0670.$$

And the probability of no more than 5 backorders is

$$\Pr(BORV_6 \leq 5) = \Pr(BpipeRV_6 \leq 35) = \sum_{k=0}^{35} \frac{e^{-34.5} (34.5)^k}{k!} = 0.578.$$

The expected backorders on day 6 with $s = 0$ (no base stock) is just

$$E[BORV(0,6)] = Bpipe(6) = 34.5.$$

Recasting Equation 3-22 to apply to expected base backorders rather than to expected depot backorders, we have the recursion formula

$$E[BORV(s+1, 6)] = E[BORV(s, 6)] - \left[1 - \sum_{k=0}^s \frac{e^{-34.5} (34.5)^k}{k!} \right].$$

Applying this formula repeatedly, we find with $s = 30$, the expected base backorders figure is $E[BORV(30, 6)] = 5.21$. Similar calculations yield the base backorder PDF and the expected backorders for days other than day 6.

BACKORDER STATISTICS WITH POSITIVE, FINITE DEPOT STOCK

Suppose we now increase the depot stock level from zero to one. We continue to focus on day 6. The means and variances of the base repair pipeline, of the order and ship pipeline, and of the depot repair pipeline on day 6 remain as before; they are not affected by the depot stock level.

We compute the mean and variance of depot backorders using the recursion relations in Equations 3-22 and 3-23, together with Equation 3-25. From Equation 3-22, we have

$$DEBO(s_0 + 1, T - OST) = DEBO(s_0, T - OST) - \left[1 - \sum_{k=0}^{s_0} P(k) \right],$$

where $P(k)$ is the probability that the $DRpipeRV = k$ at time $T - OST$. Putting $s_0 = 0$, $T = 6$, and $OST = 3$ and recalling that the depot repair pipeline at $T = 6$ is a Poisson random variable with mean 12.5, we find that the figure for expected depot backorders on day 3 with a depot stock level of 1 is

$$\begin{aligned} DEBO(1, 3) &= DEBO(0, 3) - [1 - P(0)] = DRpipe(3) - (1 - e^{-12.5}) \\ &= 12.5 - (1 - 3.7 \times 10^{-6}) \\ &= 11.5 \end{aligned}$$

(here $P(0)$ is the probability that the $DRpipe$ is zero on day 3).

Using Equation 3-23, we find that the second moment for depot backorders is

$$\begin{aligned} DE2BO(s_0 + 1, T - OST) &= DE2BO(s_0, T - OST) - DEBO(s_0, T - OST) \\ &\quad - DEBO(s_0 + 1, T - OST). \end{aligned}$$

Again putting $s_0 = 0$, $T = 6$, and $OST = 3$ and recalling Equation 3-24, we find that

$$DE2BO(0, 3) = VDBO(0, 3) + DEBO(0, 3)^2 = 12.5 + (12.5)^2 = 168.75,$$

and

$$\begin{aligned} DE2BO(1, 3) &= DE2BO(0, 3) - DEBO(0, 3) - DEBO(1, 3) \\ &= 168.75 - 12.5 - 11.5 \\ &= 144.75. \end{aligned}$$

From Equation 3-25, we find that the variance of depot backorders is

$$VDBO(1, 3) = DE2BO(1, 3) - DEBO(1, 3)^2 = 144.75 - (11.5)^2 = 12.5.$$

We now calculate the mean and variance of the base resupply pipeline on day 6 as we did before:

$$Bpipe(6) = BRpipe(6) + OSpipe(6) + DEBO(1, 3) = 14 + 8 + 11.5 = 33.5;$$

$$VBpipe(6) = VBRpipe(6) + VOSpipe(6) + VDBO(1, 3) = 14 + 8 + 12.5 = 34.5.$$

The variance-to-mean ratio of the base resupply pipeline on day 6 is $\frac{34.5}{33.5} = 1.02985$ (the calculations that follow require more precision than what we have been using). Since this figure is larger than 1, we use the negative binomial (Vari-METRIC) approximation to the PDF for the number of items in this pipeline on day 6, as given in Equation 3-15:

$$\Pr(BpipeRV_6 = k) = \frac{\Gamma\left(k + \frac{\mu(6)}{VMR(6) - 1}\right)}{k! \Gamma\left(\frac{\mu(6)}{VMR(6) - 1}\right)} \left(\frac{VMR(6) - 1}{VMR(6)}\right)^k \left(\frac{1}{VMR(6)}\right)^{\frac{\mu(6)}{VMR(6) - 1}},$$

or

$$\begin{aligned} \Pr(BpipeRV_6 = k) &= \frac{\Gamma\left(k + \frac{33.5}{1.02985 - 1}\right)}{k! \Gamma\left(\frac{33.5}{1.02985 - 1}\right)} \left(\frac{1.02985 - 1}{1.02985}\right)^k \left(\frac{1}{1.02985}\right)^{\frac{33.5}{1.02985 - 1}} \\ &= \frac{\Gamma(k + 1122.5)}{k! \Gamma(1122.5)} (0.02899)^k (0.97)^{1122.5}. \end{aligned}$$

As we did before, assume that the base stock level $s = 30$. The probability of no base backorders with $s_0 = 1$ is

$$\begin{aligned} \Pr(BORV_6 = 0) &= \Pr(BpipeRV_6 \leq 30) \\ &= \sum_{k=0}^{30} \frac{\Gamma(k + 1122.5)}{k! \Gamma(1122.5)} (0.02899)^k (0.97)^{1122.5} \\ &= 0.313; \end{aligned}$$

for comparison, the probability of no backorders at the base with no depot stock was 0.253. (Note: Because of the difficulty in computing ratios of large values of the gamma function directly, we have computed this probability using the recursion formula for the negative binomial distribution — see Appendix C).

The probability of exactly 5 base backorders with $s = 30$ and $s_0 = 1$ is

$$\begin{aligned}\Pr(BORV_6 = 5) &= \Pr(BpipeRV_6 = 35) \\ &= \frac{\Gamma(35 + 1122.5)}{35! \Gamma(1122.5)} (0.02899)^{35} (0.97)^{1122.5} \\ &= 0.0642,\end{aligned}$$

and the probability of no more than 5 base backorders is

$$\begin{aligned}\Pr(BORV_6 \leq 5) &= \Pr(BpipeRV_6 \leq 35) \\ &= \sum_{k=0}^{35} \frac{\Gamma(k + 1122.5)}{k! \Gamma(1122.5)} (0.02899)^k (0.97)^{1122.5} \\ &= 0.643.\end{aligned}$$

For comparison, these probabilities were 0.067 and 0.578, respectively, with no depot stock. These comparisons illustrate the fact that adding depot stock has reduced the likelihood of any particular number of base backorders.

The expected backorders on day 6 with $s_0 = 0$ (no base stock) is just $E[BORV(0, 6)] = Bpipe(6) = 33.5$. Using the recursion formula for expected backorders, we have

$$E[BORV(s + 1, 6)] = E[BORV(s, 6)] - 1 - \sum_{k=0}^s \frac{\Gamma(k + 1122.5)}{k! \Gamma(1122.5)} (0.02899)^k (0.97)^{1122.5}.$$

Applying this repeatedly, we find that with $s = 30$ and $s_0 = 1$, the expected base backorders figure is 4.46. For comparison, the expected base backorders figure with $s = 30$ and $s_0 = 0$ was 5.21, so we see that, in this example, increasing the depot stock level by 1 has brought down the mean number of backorders on day 6 significantly.

SUMMARY

In this appendix, we have shown how to compute the means and variances of the various components of the base resupply pipeline, and how to use these pipeline statistics to approximate the PDF and the mean of base backorders. We have done this under the assumption that flying hours, and therefore the demand process, may change over time, but that the resupply times are constant. In the case of varying resupply times, or in the case where we consider backorders of lower indenture items, our approach is quite similar.

APPENDIX C

Recursion Formula for Negative Binomial Distribution

Recursion Formula for Negative Binomial Distribution

Appendix B discusses methods for computing the reduction in expected backorders resulting from adding spares. That appendix used a recursion formula for the negative binomial distribution, which we present below.

The value of the negative binomial distribution with parameters n and p at a positive integer x is given by

$$N_b(x|n, p) = \frac{\Gamma(x+n)}{x!\Gamma(n)} p^n (1-p)^x,$$

where Γ denotes the gamma function, defined for real $u > 0$ by

$$\Gamma(u) = \int_0^{\infty} t^{u-1} e^{-t} dt.$$

It can be shown that the negative binomial distribution has mean $\mu = \frac{n(1-p)}{p}$ and variance $\sigma^2 = \frac{n(1-p)}{p^2}$. (For example, a negative binomial random variable can be expressed as a sum of independent geometric random variables, whose means and variances are easy to compute.¹ In Chapter 3, we know μ and σ^2 , rather than n and p , but straightforward algebra shows that $p = \frac{\mu}{\sigma^2}$, and $n = \frac{\mu}{\frac{\sigma^2}{\mu} - 1}$.

Using the fact that for any integer, $x \geq 1$, $\Gamma(x+1) = x\Gamma(x)$, we see that for all such x ,

$$\begin{aligned} N_b(x|n, p) &= \frac{\Gamma(x+n)}{x!\Gamma(n)} p^n (1-p)^x \\ &= \left[\frac{x+n-1}{x} (1-p) \right] \left[\frac{\Gamma(x+n-1)}{(x-1)!\Gamma(n)} p^n (1-p)^{x-1} \right] \\ &= \left[\frac{x+n-1}{x} (1-p) \right] N_b(x-1|n, p). \end{aligned}$$

¹William Feller, *An Introduction to Probability Theory and Its Applications*, New York: John Wiley and Sons, Inc., Vol. 1, 3rd Edition, 1968.

APPENDIX D

Glossary

Glossary

AAM	=	Aircraft Availability Model
AOG	=	aircraft on ground
AP	=	application percentage
ASM	=	Aircraft Sustainability Model
AWP	=	awaiting parts
BO	=	backorders
BORV	=	number of base backorders (random variable)
Bpipe	=	base resupply pipeline (physical process) or mean number of items in that pipeline
BpipeRV	=	number of units in base resupply pipeline (random variable)
BRpipe	=	base repair pipeline (physical process) or mean number of items in that pipeline
BRpipeRV	=	number of units in base repair pipeline (random variable)
BRT	=	base repair time
BRTNHA	=	base repair time for next higher assembly
BRTLRLU	=	base repair time for LRU
CDF	=	cumulative distribution function
ConPCT	=	condemnation percentage
DBORV	=	number of depot backorders (random variable)
DEBO	=	depot expected backorders
DoD	=	Department of Defense
DI	=	due in

DO	=	due out
DRT	=	depot repair time
DRpipe	=	depot repair pipeline (physical process) or mean number of items in that pipeline
DRpipeRV	=	number of units in depot repair pipeline (random variable)
DSO	=	direct support objective
EBO	=	expected backorders
ENMCS	=	expected not mission capable – supply
FAP	=	future application percentage
FF	=	failure factor (failures per flying hour)
FHP	=	flying-hour program
FIFO	=	first in, first out
FIT	=	fault isolation time
IAF	=	Israel Air Force
ISAAC	=	Initial Spares Aircraft Availability Calculation
$I(t)$	=	induction time for item emerging from pipeline at time t
JSTARS	=	Joint Surveillance Target Attack Radar System
LMI	=	Logistics Management Institute
LRU	=	line replaceable unit
METRIC	=	Multi-Echelon Technique for Recoverable Items Control
MRSP	=	mobility readiness spares package
NAC	=	number of aircraft
NHA	=	next higher assembly
NMCB	=	not mission capable – both (maintenance and supply)
NMCM	=	not mission capable – maintenance

NMCS	=	not mission capable – supply
NRTS	=	not reparable this station
OH	=	on hand
OSpipe	=	order and ship pipeline (physical process) or mean number of items in that pipeline
OSpipeRV	=	number of units in order and ship pipeline (random variable)
OST	=	order and ship time
PDF	=	probability distribution function
PLT	=	procurement lead-time
PLTT	=	procurement lead-time total
Q	=	order quantity
QPA	=	quantity per application
r	=	reorder point
RAT	=	reassembly time
RR	=	mean repair rate
RT	=	mean repair time
SRU	=	shop replaceable unit
ST	=	resupply suspension time
TI _{<i>l</i>}	=	total installed for part <i>l</i>
TNMCS	=	total not mission capable – supply
TOIMDR	=	total organizational and intermediate maintenance demand rate
USAF	=	United States Air Force
Vari-METRIC	=	model based on METRIC model with better treatment of pipeline variance

VBRpipe = variance of base repair pipeline

VDBO = variance of depot backorders

VMR = variance-to-mean ratio

REPORT DOCUMENTATION PAGE

Form Approved
OPM No.0704-0188

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 reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Information and Regulatory Affairs, Office of Management and Budget, Washington, DC 20503.

1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE Oct 96		3. REPORT TYPE AND DATES COVERED Final	
4. TITLE AND SUBTITLE Optimizing Spares Support: The Aircraft sustainability Model				5. FUNDING NUMBERS C DASW01-95-C-0019 PE 0902198D	
6. AUTHOR(S) F. Michael Slay, Tovey C. Bachman, Robert C. Kline, T. J. O'Malley, Frank L. Eichorn, Randall M. King					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Logistics Management Institute 2000 Corporate Ridge McLean, VA 22102-7805				8. PERFORMING ORGANIZATION REPORT NUMBER LMI- AF501MR1	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) HQ USAF/LG Room 4E260, The Pentagon Washington, DC 20330				10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES					
12a. DISTRIBUTION/AVAILABILITY STATEMENT A: Approved for public release; distribution unlimited				12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) Modern inventory management typically relies, when possible, on quick deliveries from suppliers, and aims to carry minimal inventory. Such practices work best in managing low-cost items with stable demand and significant market size. But supporting advanced military equipment poses unique problems: small fleet sizes and low operating tempos lead to sporadic and unstable demand patterns; components are highly specialized, have limited or uncertain sources of supply, and often procurement lead-times of years. The Aircraft Sustainability Model (ASM), developed by the Logistics Management Institute for the United States Air Force, is a mathematical model that computes optimal spares mixes to support a wide range of possible operating scenarios. The ASM sizes spares levels based on desired weapon system readiness levels, rather than on supply-oriented measures, such as stock on shelf or percent of demands filled. The ASM has been used by the USAF to determine spares kits to support wartime squadron deployments; an enhanced version handles initial provisioning. This report describes the model, the problems it was developed to solve, and the mathematical techniques it uses to solve them. It is intended for readers with some background in probability and statistics. Familiarity with multi-echelon inventory theory is not required.					
14. SUBJECT TERMS Aircraft, availability, cannibalization, deployment, initial provisioning, inventory, inventory management, inventory model, multi-echelon, multi-indenture, optimization, readiness, resource allocation, safety level, spares, stock level, stockage policy, supply, supply management, supply model, sustainability, war reserve, weapon system support				15. NUMBER OF PAGES 154	
				16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL		