

19961114 063

RAND

*Government Contracting
Options*

A Model and Application

Edward G. Keating

50th
Project AIR FORCE
1946-1996

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

DTIC QUALITY INSPECTED 1

The research reported here was sponsored by the United States Air Force under Contract F49642-96-C-0001. Further information may be obtained from the Strategic Planning Division, Directorate of Plans, Hq USAF.

Library of Congress Cataloging in Publication Data

Keating, Edward G. (Edward Geoffrey), 1965–
Government contracting options : a model and application /
Edward G. Keating.
p. cm.
"Project AIR FORCE."
"Prepared for the United States Air Force."
"MR-693-AF."
Includes bibliographical references.
ISBN 0-8330-2422-1 (alk. paper)
1. Defense contracts—United States. 2. United States. Air
Force—Equipment—Maintenance and repair. I. RAND
Corporation. II. Project AIR FORCE (U.S.) III. United
States. Air Force. IV. Title.
UG1123.K43 1996
358.4'16211'0973—dc20

96-28078
CIP

© Copyright 1996 RAND

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from RAND.

RAND is a nonprofit institution that helps improve public policy through research and analysis. RAND's publications do not necessarily reflect the opinions or policies of its research sponsors.

Published 1996 by RAND
1700 Main Street, P.O. Box 2138, Santa Monica, CA 90407-2138
RAND URL: <http://www.rand.org/>

To order RAND documents or to obtain additional information, contact Distribution Services: Telephone: (310) 451-7002; Fax: (310) 451-6915; Internet: order@rand.org

RAND

*Government Contracting
Options*

A Model and Application

Edward G. Keating

*Prepared for the
United States Air Force*

50th
Project AIR FORCE
1946-1996

DTIC CURRENT INFORMATION

Approved for public release; distribution unlimited

PREFACE

This report considers how Air Force repair contracts should be designed. A large and potentially growing fraction of repairs are done by contractors. Can contracts be redesigned to provide contractors with better incentives to have a robust, high-speed repair system?

This report models and simulates contractor decision making. Insights are then derived as to how repair contracts should be designed to induce optimal contractor behavior.

This work was carried out in the Contracting for Lean Logistics study in the Resource Management Program of Project AIR FORCE, RAND's federally funded research and development center (FFRDC) funded by the United States Air Force. It was sponsored by the Deputy Chief of Staff/Logistics, Headquarters, USAF. It should be of interest to logistics managers and analysts throughout the Air Force logistics system, especially those involved in contracting, and to logisticians and contracting officials in the other military departments and in the Office of the Secretary of Defense.

PROJECT AIR FORCE

Project AIR FORCE, a division of RAND, is the Air Force federally funded research and development center (FFRDC) for studies and analyses. It provides the Air Force with independent analyses of policy alternatives affecting the development, employment, combat readiness, and support of current and future aerospace forces. Research is being performed in three programs: Strategy and Doctrine, Force Modernization and Employment, and Resource Management and System Acquisition.

In 1996, Project AIR FORCE is celebrating 50 years of service to the United States Air Force. Project AIR FORCE began in March 1946 as Project RAND at Douglas Aircraft Company, under contract to the Army Air Forces. Two years later, the project became the foundation of a new, private nonprofit institution to improve public policy through research and analysis for the public welfare and security of the United States—what is known today as RAND.

CONTENTS

Preface..... iii

Figures..... vii

Tables..... ix

Summary..... xi

Acknowledgments..... xv

1. INTRODUCTION 1

 Background..... 1

 A Model..... 2

 Simulation and Nonlinear Search 4

2. AN APPLICATION AND SIMULATION 6

 No Per-Unit Penalty?..... 14

 The Cost-Availability Curve 15

 Optimality and Robustness 18

3. EFFECTS OF CHANGING ASSUMPTIONS 19

 A Risk-Neutral Contractor 19

 Government Holds Incorrect Distributional Beliefs 22

 Prohibitively Expensive Modifications 24

 Hands-On Repair Speed Is Costly 25

 Spares Half as Costly..... 27

 Different Contracting Approaches 28

 One-Year Contract 28

 Contractor Responsible for Ten Times Fewer Aircraft 29

4. CONCLUSION AND IMPLICATIONS 31

Appendix

A. BASIC MODEL AND ENHANCEMENTS 33

B. THE CONTRACT SIMULATION PROGRAM 43

C. APPLICATION PARAMETERS 45

References..... 49

FIGURES

1. Sketch of the Model	4
2. Risk-Averse Contractor Utility Function	8
3. Risk-Averse Contractor Profits with Optimal and Lump-Sum-Only Contracts	13
4. Cost-Availability Curve	17
5. Risk-Neutral Contractor Profits with Lump-Sum-Only Contract	21
6. Basic Model	35
7. Unit Cost of a Repair	47

TABLES

1. Base Case Application Parameters	7
2. Simulation of Time and Materials Contract with Risk-Averse Contractor, 98% Required Aircraft Availability	10
3. Simulation Contract Approaches with Risk-Averse Contractor, 98% Required Aircraft Availability	12
4. Simulation Contract Approaches with Risk-Averse Contractor, 98% Required Aircraft Availability	15
5. Optimal Contracts and Required Aircraft Availability Levels .	16
6. Simulation Contract Approaches with Risk-Neutral Contractor, 98% Required Aircraft Availability	20
7. Simulation Contract Approaches with Risk-Averse Contractor, Failure Fraction Half Expected Level, 98% Required Aircraft Availability.....	22
8. Simulation Contract Approaches with Risk-Averse Contractor, Failure Fraction Twice Expected Level, 98% Required Aircraft Availability.....	23
9. Simulation Contract Approaches with Risk-Averse Contractor, Prohibitively Expensive Modifications, 98% Required Aircraft Availability.....	24
10. Simulation Contract Approaches with Risk-Averse Contractor, Hands-On Repair Speed Costs \$500,000 per Day, 98% Required Aircraft Availability	26
11. Simulation Contract Approaches with Risk-Averse Contractor, Hands-On Repair Speed Costs \$1,500,000 per Day, 98% Required Aircraft Availability	26
12. Simulation Contract Approaches with Risk-Averse Contractor, Spares Cost Half as Much, 98% Required Aircraft Availability	27
13. Simulation Contract Approaches with Risk-Averse Contractor, One-Year Contract, 98% Required Aircraft Availability	28
14. Simulation Contract Approaches with Risk-Averse Contractor, Contractor Responsible for Ten Times Fewer Aircraft, 98% Required Aircraft Availability	29

SUMMARY

Contractors represent a sizable, and potentially growing, portion of the Air Force's repair system. How should the Air Force design its repair contracts?

In this report, we develop an economic model of contractor motivations and behavior and simulate how contractors would respond to different types of contracts. We derive the government's optimal contract under varying scenarios.

This model of contractor behavior is useful as a way to quickly and inexpensively test different types of contracts. Further, models are a well-developed and accepted part of economic theory,¹ and that research is utilized here in the construction and interpretation of the model.

A MODEL

We model an aircraft system that experiences stochastic failures. Broken parts enter the repair system irregularly. The contractor must repair the broken parts and/or replace them with spare parts to maintain a specified aircraft availability level.

The contractor has a variety of choice variables, e.g., repair capacity and quality. We assume the contractor makes these choices based on what course of action will prove best for the contractor, given the contract provided by the government. Meanwhile, the government chooses the contract form, which may include stipulations regarding a fee per unit repaired, a fee per spare required, and/or a lump-sum fee that does not vary with the number of units repaired or spares needed. The government knows the contractor will maximize for its own benefit in response to the contract provided. In the model, the government must provide the contractor with a combination of fees and a lump sum that is lucrative enough *ex ante* (ahead of time) to induce the contractor to participate in the contract.

¹See, for example, Schmalensee and Willig (1989).

AN APPLICATION AND SIMULATION

We used data based on select F-16 components to run the model and understand contractors' and government behavior. Our simulations suggest that the type of contract that the government provides has a large impact on contractor behavior. For example, a contract that pays a contractor a sizable fee per unit repaired produces very poor incentives in this illustrative analysis. The contractor's incentive is to choose very low-quality repair to assure itself of a steady stream of broken parts in the future. The government's expected expenditures are quite large with the contracts of this sort that we analyzed. Unfortunately, this is a very common contracting approach.

Our simulation suggests a contract that combines a sizable lump-sum payment with cost-sharing for required expensive spares can be a desirable approach. With such a contract, the contractor receives a lump-sum payment that does not vary with the number of units repaired. The contractor also receives partial reimbursement if expensive spares are needed. In exchange, the contractor agrees to maintain a specified aircraft availability rate. Contractors compete for the contract on the basis of who will accept the lowest lump-sum payment, holding fixed the required availability level and the spares cost-sharing formula. The model indicates that it is better to reward a contractor for meeting an aircraft availability target rather than compensating the contractor per item repaired. An availability-oriented contract encourages high-quality contractor repair; the contractor benefits from high-quality repair since fewer broken items appear subsequently. A contract with a lump-sum and expensive item cost-sharing also encourages rapid contractor repair, even without any explicit rewards for fast repair turnaround.

EFFECTS OF CHANGING ASSUMPTIONS

We derived an optimal contract for a specific set of parameters. To test the robustness of this result, we analyzed the performance of this contract in case various parameters varied.

In general, the base-case optimal contracts performed well in the face of parameter perturbations. For example, although the optimal

contract with a risk-neutral contractor involves no spares cost-sharing, the contract that was optimal with a highly risk-averse contractor still performs adequately with contractor risk-neutrality. A halving or doubling of the part failure rate relative to the government's expectations is virtually irrelevant provided the contractor knows the actual failure rates. As one might expect, large-scale changes in the contractor's repair cost function change the government's expected expenditures, but the form of the contract does not change in important ways. One clear source of difficulty, though, would be if the government overestimates the costs of new spares and offers the contractor more per new spare required than the spares cost the contractor. Also, we find the optimal contract does vary in important ways if contract duration is changed.

CONCLUSION AND IMPLICATIONS

Our model and simulations suggest that a contract with a lump-sum and expensive item cost-sharing in which payment to the contractor does not vary with the number of units repaired deserves greater attention in government repair contracting. When the government can rely on contractor competition to arrive at the lump-sum payment, our simulations suggest such contracts place very limited informational requirements on the government. The government does not need vast knowledge of part failure rates or contractor costs. Such a contract does, however, assume the contractor has fairly detailed information about a weapon system, e.g., its failure pattern. Such contracts are probably most appropriate, therefore, for mature weapon systems with predictable usage patterns.

ACKNOWLEDGMENTS

The author has benefited greatly from the guidance, comments, and assistance of numerous people. Ray Pyles was the project leader on this effort. Frank Camm was the program director when this endeavor commenced. Laura Baldwin, Mary Chenoweth, Shirley Cromb, Chris Hanks, Jan Harris, Jim Hosek, Penny Koerner, Craig Moore, Adele Palmer, Don Palmer, Ellen Pint, Ken Reynolds, Bob Roll, Hy Shulman, Dennis Smallwood, and Jerry Sollinger assisted with this document. Tony Bower and Ken Girardini thoughtfully and carefully reviewed this document. Also, Laura Zakaras lent her considerable skills as a communications analyst. Linda Lenhoff skillfully edited this document. Phillip Feldman, Jan Hanley, Tim Ramey, and Terry West helped with the computer simulation. Lionel Galway and Tom Lucas provided statistical advice. Susan McGlamery provided valuable library research support. Jonathan Feinstein was a coauthor on an earlier paper that presaged many of this report's approaches; Jonathan also provided helpful comments on this report. This work was presented at the Southern Economics Association conference in New Orleans on November 18, 1995. The author thanks seminar participants and discussants, especially Barbara Alexander and Thomas Zak. The author also thanks participants at seminars at RAND on August 5, 1994, August 23, 1994, and June 23, 1995. Of course, remaining errors are the author's responsibility.

1. INTRODUCTION

BACKGROUND

In Fiscal Year 1995, the Department of Defense spent in excess of \$13 billion on depot-level weapon system maintenance and repair. The Air Force's share of this spending exceeded \$4 billion.¹ The Air Force spent in excess of \$1 billion on private sector contractor repair.² Further, proposals have been made to eliminate or lessen the current legislative mandate that at least 60 percent of maintenance work must go to public sector depots.³ Hence, contractor repair and maintenance is a large spending category and it may be growing further. Clearly, it is important and valuable that the government design these repair contracts as carefully and thoughtfully as possible.

In light of the large and potentially growing role contractors play in these processes, how might repair contractor performance be improved? How can component repair contracts be modified to ensure that contractors will take effective action to minimize their flow times and enhance the effectiveness of their efforts? Can contractors be induced to provide high-quality, responsive repair?

The United States government has typically used "per-repair" contracts. With a per-repair contract, the contractor receives compensation that is a direct function of the number of items the contractor repairs.

One version of a per-repair contract is a requirements contract in which the contractor receives a specified fee per unit repair without a guaranteed minimum workload. Time and material contracts are also a version of a per-repair contract. With a time and materials contract, the contractor is paid for whatever labor time and materials are

¹See United States Air Force, Defense Business Operations Fund. FY 1995 Budget Estimates. U.S. Air Force Overview. Operating/Capital Budget (1994).

²Leland (1995) notes that about 70 percent of the Air Force's total maintenance is done at public depots; the rest goes to contractors.

³See Commission on Roles and Missions of the Armed Forces (1995) and Leland (1995).

required to fix broken items. Also in the category of per-repair contracts, the government sometimes uses fixed price contracts where a fixed number of items are guaranteed to be entering the repair process.

Per-repair contracts have been criticized for not providing contractors with good incentives. For example, because the contractor gets paid additionally each time a piece of equipment needs to be repaired, the contractor lacks obvious incentive to do high-quality repair.

To address this problem, the government has begun to use a variety of other contracting approaches that move beyond per-repair contracts. One approach we label an "availability-oriented" contract; i.e., the contractor is paid for keeping a system operating and available. For example, Serv-Air Inc. has a contract with the Air Force for repair of the C-21. The C-21 is the military version of a Lear jet. Serv-Air has a ten-year contract with the Air Force for C-21 repair and is compensated on the basis of the number of C-21 flight hours the Air Force receives. If, for example, Serv-Air can increase C-21 reliability through improved repair processes, they will benefit financially even if fewer C-21 components then enter the repair process.

Similarly, the Navy has a five-year contract with Litton Industries for repair of the LN-15C Inertial Navigation Unit. Litton receives a payment per LN-15C flight hour while guaranteeing the Navy a specified availability level. As with the Serv-Air contract, this LN-15C contractual arrangement provides Litton with enhanced incentives to make repair more lasting and effective. In contrast, improved repair processes can reduce short-run firm revenue with a per-repair contract.

Our objective in this research was to model repair contractor behavior to gain general insights into appropriate contract formulation. What incentives do per-repair contracts provide? Does the C-21 or the LN-15C approach seem more reasonable? Are there other desirable options?

A MODEL

This report describes a model that the Air Force (or any other organization interested in outsourcing its maintenance function) could

use to help assess and compare different ways of designing the contract for maintenance services. It provides a flexible way to test the efficacy of different types of repair contracts. As currently constructed, the model allows a risk-averse contractor to choose repair capacity, repair quality, the modification level, and hands-on repair time so as to maximize its utility, given the contractual terms that the government offers. The model then uses simulation and numerical parameter search to predict how a utility-maximizing contractor would respond to any government offer and uses such predictions to derive the optimal contract from the government's point of view.

The model can accommodate a wide range of assumptions about the characteristics of the contractor, the production process, the state of the world, and the contract designs to be evaluated. We used a specific example based on Air Force maintenance data to illustrate the range of questions that can be addressed and to draw contracting policy implications.

Figure 1 summarizes the model we developed. In the model, there is a flight line that generates broken parts. These parts flow to the contractor who is supposed to repair them.

The repair contractor has a variety of choices to make. For example, the contractor makes a one-time choice of how much repair capacity to build, how much quality to put into repairs, and how much time the hands-on repair takes. Appendix A explains these parameters and details the mathematics of the model.

The contractor's preferences will influence its behavior. One view is that contractors are risk-neutral profit-maximizers. However, another view is that contractors may be risk-averse. Risk-averse contractors put comparatively greater negative weight on possible losses of a certain level than they put positive weight on possible profits of that same level. We examine both these cases using the model.

Meanwhile, the government has to decide what sort of contract to give to the contractor. The government could set an aircraft availability requirement. The government could give the contractor a fixed ("lump-sum") payment that does not vary with the number of items repaired or spares required. The government could also give the

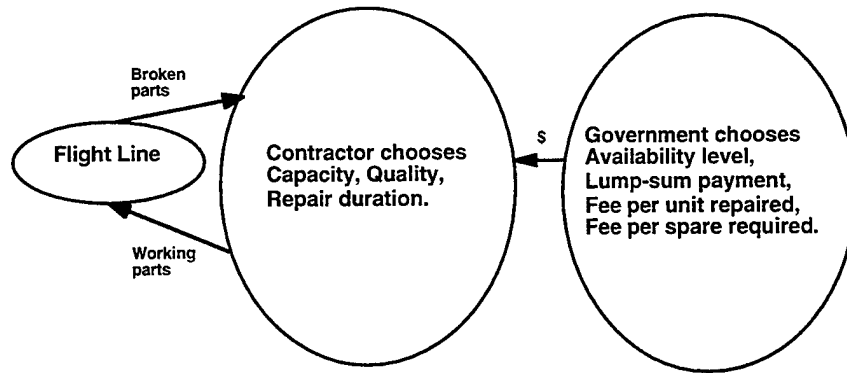


Figure 1-Sketch of the Model

contractor funds every time a repair is done or a spare is needed. The different approaches to contracting discussed earlier are subsumed in these cases. For example, a requirements contract lacks an availability requirement or a lump-sum payment, but has a fee per unit repaired. The C-21 and LN-15C contracts, in contrast, have availability requirements and lump-sum payments.

SIMULATION AND NONLINEAR SEARCH

Given this model, we then simulate it. We use sets of random numbers to simulate streams of broken aircraft. We then use nonlinear parameter search to compute how the contractor best responds to possible broken aircraft streams, e.g., what level of repair capacity and quality is best for the contractor. Finally, we use another nonlinear search to compute the best contract specifications for the government to provide to the contractor, given the government knows the contractor will respond in whatever manner serves its ends in response to that contract.

It is important to understand the limitations of any simulation process of the sort used in this report. For example, any simulation inherently involves a finite number of random draws from asserted statistical distributions. We believe we have undertaken enough draws (generally 50 sets of 1250 periods) to find stable results. However, one can never be fully sure. For example, changes in the random number generator seed (building off a different string of random numbers) will result in somewhat different results. We believe, however, that the basic findings, as opposed to the specific numeric results, in this report are robust to such changes.

Further, our simulations' searches for the contractor's choices and the government's optimal contract involve nonlinear optimization. We use a nonlinear search algorithm⁴ that we believe to be robust and stable. However, nonlinear search algorithms involve starting values (i.e., best guesses of where solutions might lie), and it is always possible that nonlinear search algorithm solutions will vary with the starting values provided. Again, our experiments with different starting values in our program lead us to believe our findings are robust. Further, the contractor optima are similar case to case. But the nature of nonlinear search is that there can be no guarantees that global optima have been found in every case.

In Chapter 2, we simulate this model using the Radio Frequency (RF) repair stand and its repairs of F-16 transmitters and antennae as an example. We derive the optimal contract given our parameters. Chapter 3 examines the sensitivity of the results to variations in our assumptions. We show that Chapter 2's optimal contract continues to perform fairly well even if the government's knowledge of the real world parameters is imperfect. Chapter 4 concludes the report.

⁴We use the AMOEBA search algorithm described by Press et al. (1989) on pages 326-330. Appendix B discusses our search algorithm in more depth.

2. AN APPLICATION AND SIMULATION

We next apply this model to a realistic problem, the RF repair stand and its repairs of F-16 transmitters and antennae, and simulate the model. The primary goal in this chapter is to derive the "optimal contract" in this example, i.e., the contract that minimizes expected government expenditure given that the risk-averse contractor maximizes in response to the contract it gets. We also test the efficacy of common contracting approaches, e.g., time-and-materials contracts. We assume a structure in which the government declares what per-unit repaired fee and spares cost-sharing it is offering, along with the required aircraft availability rate, and contractors compete based on who will accept the lowest lump-sum payment. Later in this chapter, we show how the optimal contract changes as the required aircraft availability level changes. In Chapter 3, we show how contractor behavior and the optimal contract change when model parameters change.

In this example, we assume one has 403 F-16 aircraft that must be kept at a 98 percent Fully Mission Capable (FMC) rate at all times, i.e., 395 aircraft available. We assume an available aircraft averages 0.9 sorties per day, five days per week, with 1.25 hours per sortie. We are using a specific type of contract in which the contractor must guarantee a specified availability level, in this case 98 percent.

This F-16 example is based on the Air Force's Coronet Deuce exercise described in Abell and Shulman (1992). The parameters we assume are designed to add realism to this example but are not to be viewed as real or valid. The parameters chosen are important, however, in that they drive the specific contract parameters we derive.

Table 1 and Appendix C present our base-case parameter assumptions.

We assume the contractor is risk-averse and has the utility function

$$U(k) = k^{\frac{1}{2}}, k \geq 0$$
$$U(k) = -k^2, k \leq 0$$

where k is firm profit or losses, measured in thousands of dollars. Figure 2 depicts this utility function.

Table 1
Base Case Application Parameters

Parameter	Value
Cost per new part	\$198,578
Statistical distribution of part failure fraction	Lognormal with mean 0.013, variance twice as large as the mean, truncated at 1.
Annual cost per repair stand	\$535,000.00
Periods per year	250
Statistical distribution of repair stand failure fraction	Lognormal with mean 0.1 and variance 0.2, truncated at 1.
Repair quality per-unit cost function	$\$6407.8 * q^{1.68}$
Modification level per-unit cost function	$\$39715.70 * m^2 / (1 - m)$
Modification level fixed cost function	$\$10,000,000 * m$
Cost of speeding hands-on repair time by one day	0

NOTE: See Appendix C for details.

This contractor is quite risk-averse. Contractor utility could be defined in other ways, and/or different levels of risk-aversion could be chosen. A general formulation is that the contractor's utility function is $U(k(x,q,m,t_r))$, where x is contractor repair capacity, q is contractor repair quality, m is the contractor's modification level, and t_r is the contractor's choice for the hands-on repair time.

In this application, the government has three choice variables. The first is the fee per unit repaired ($c1$). The second is the fee per new unit required ($c2$). Finally, the government can make a lump-sum payment to the contractor. In this application, we do not consider fees based directly on the contractor's costs. Instead, the fees are based on the number of units repaired and new units needed. We wished to




Figure 2--Risk-Averse Contractor Utility Function

avoid the real-world complexity of government monitoring of contractor costs.

Further, the lump-sum payment is constrained as the government must be sure the contractor's expected utility is at least zero. The government cannot compel a contractor to perform; a contractor must be offered arrangements whereby its expected utility is nonnegative. If we let E denote the government's expected expenditure, the government wishes to choose $c1$ and $c2$ to minimize $E(c1,c2)$.

In response to any contract, the contractor has four choices: the number of repair stands to buy (x - fractional values are allowable), the quality of repairs (q falls between 0 and 1), the modification level to choose (m falls between 0 and 1), and the hands-on repair time (t_r - must be an integer). We assume the contractor knows everything in Table 1, including the statistical distributions of part and repair station failure.

To summarize, we have the government choosing $c1$ and $c2$ to minimize $E(c1, c2)$ subject to the contractor choosing x , q , m , and t_r to maximize $U(k(x, q, m, t_r))$.

We simulate this model over five 250-period "years". (There are approximately 250 business days in a calendar year.) Appendix B describes the simulation procedure. We assume a repair station can be rented for five years for \$2.675 million. We do not consider the possibility of contractual breach by either the contractor or the government.

Before finding the optimal $c1$ and $c2$, we examined a "time and materials" contract. With this type of contract, the government gives the contractor a fee equal to the marginal cost of $q=1$ repair every time an item is repaired. Also, the government promises to pay for any spares the contractor needs. The government also offers the contractor a lump-sum payment to make the contractor's expected utility zero. Table 2 shows how the contractor responds to such a contract.

The contractor chooses a high-capacity level, but zero quality and no modifications. The coexistence of a high-capacity level with low quality and no modifications occurs because the contractor is paid a high $c1$ of \$6,408 every time a broken item enters the repair process. Hence, in this case, high capacity is chosen to handle high volumes of repair that generate a great deal of fee income. However, once in the repair process, the contractor has no incentive to undertake high-quality repair. Instead, the contractor wants items to break again. The contractor relies on the government-funded spares ($c2 = 198,578.47$ or 100 percent of the spares cost) to maintain the required 98 percent aircraft availability level.

Table 2
Simulation of Time and Materials Contract with Risk-Averse Contractor
98% Required Aircraft Availability

	Parameter	Time and Materials Contract
Government	$c1$	\$6,407.80
	$c2$	\$198,578.47
	Lump-sum	-\$1,598,990,000
Contractor	x	416.13
	q	0.00
	m	0.00
	t_r	1
	Government	Expected expenditure

NOTE: Simulation based on 50 sets of 1250 periods. Expected contractor utility is zero for each contract.

Table 2's results are disturbing. Table 2 suggests that a contractor has a strong incentive to cheat on a "time and materials" contract. Instead of choosing high-quality ($q=1$) repair, the contractor chooses low-quality ($q=0$) repair. Broken items then churn through the contractor's repair facility with the contractor receiving the \$6,407.80 fee for each one. A fee of \$6,407.80 per unit "repaired" does not, in this case, induce the contractor to choose the quality level associated with a contractor marginal cost of \$6,407.80. Instead, $q=0$ and $m=0$ repair with no marginal cost to the contractor is chosen and each $c1=6407.80$ temporarily becomes contractor profit. Ultimately, however, this profit is competed away in the lump-sum.

Interfirm competition does not move the government away from Table 2's bad outcome. Indeed, the contractor's expected utility in Table 2 is zero. The large negative lump sum implies the time and materials contract is not ultimately lucrative to the contractor.¹ However, Table 2's contract is extremely costly to the government.

¹Of course, many actual defense contracts do have a separate fee based on the size of the contract. Tying the contractor's profit to the size of the contract will clearly impact the contractor's incentives.

This is a noteworthy result in that the government often uses time and materials contracting approaches. There is an allure to a contract in which the government pays for exactly the repairs that occur. Table 2's result suggests, however, that there can be severe incentive problems emanating from this sort of arrangement. Contractors are tempted to skimp on repairs to the extent they can fool the government into believing expensive repair has occurred.

Clearly, the legal system, reputational effects, monitoring of the contractor, and other mechanisms can be used to mitigate this bad outcome. However, left unchecked, the model shows that time and materials contracts have a tendency to induce low quality.

However, it is noteworthy that, in this simple, illustrative case, "good" outcomes can be elicited without intensive monitoring of contractor quality. Table 3 shows that the government can do considerably better than offering Table 2's contract. Table 3's optimal contract was derived through a nonlinear search of potential contracts the government might offer. In other words, given that we know the contractor will optimize in response to any government contract, we found the Table 3 optimal contract provides the best expected results from the government's perspective. We have one nonlinear search in which the contractor finds its best response to the government's contract and another search in which the government finds its best contract, knowing the contractor will optimize with respect to that contract.

In Table 3's optimal contract, there is a penalty for having to repair a part (c_1), but now the government pays 88 percent ($\$174,110/\$198,578.47$) of the cost of buying a new part. The risk-averse contractor wants protection against needing expensive spares. However, subsidizing spares lessens the contractor's incentive to choose high quality. Hence, a per-repair penalty ($c_1 < 0$) accompanies spare cost-sharing.

The contractor receives a big lump-sum payment from the government. (This lump-sum payment is chosen to set expected contractor utility equal to zero given c_1 and c_2 . We envision contractors competing on the basis of who will accept the lowest lump-sum payment, holding fixed

Table 3
Simulation Contract Approaches with Risk-Averse Contractor
98% Required Aircraft Availability

	Parameter	Optimal Contract	Lump-Sum-Only Contract
Government	<i>c1</i>	-\$2,628	\$0
	<i>c2</i>	\$174,110	\$0
	Lump-sum	\$73,292,700	\$122,728,000
Contractor	<i>x</i>	5.13	11.06
	<i>q</i>	1.00	1.00
	<i>m</i>	0.53	0.67
	<i>t_r</i>	1	1
	Expected expenditure	\$101,025,900	\$122,728,000

NOTE: Simulation based on 50 sets of 1250 periods. Expected contractor utility is zero under each contract.

c1 and *c2*. An alternative approach would be to have the lump-sum provide the contractor with a fixed profit level.) This lump-sum payment could be spread over time to reduce the possibility of contractor default. The important point, however, is that the lump-sum payment does not vary with the number of broken parts or spares required.

A fast hands-on repair time is chosen even though no explicit bonus is paid for speed of repair. It is possible to have contracts with fast hands-on repair time even without explicitly rewarding or requiring speed. When the contractor knows it will have to buy expensive spares if availability problems occur, it is in the contractor's best interests to move broken items through the repair process rapidly.

The average repair capacity (*x*) utilization rate associated with Table 3's optimal contract is approximately 44 percent. However, the repair capacity is heavily used early in the contract. With a positive *m*, though, parts do not fail as often after being repaired once. As time passes, a growing percentage of the parts has been modified. Hence, the number of failures and repair stand usage decrease over the duration of the contract.

In Table 3, the pure lump-sum contract is not optimal and its cost is about 21 percent greater than the cost of the optimal contract. The

government incurs additional cost with lump-sum contracts due to contractor risk-aversion. The risk-averse contractor chooses a somewhat higher modification level (m) and considerably greater capacity level (x) than is optimal. The government ultimately pays for this increased capacity and modification level through the lump-sum payment.

Figure 3 shows the frequency of different ranges of realized contractor profits from 50 sets of simulations of Table 3's optimal contract and the lump-sum-only contract. In all but one set of draws, the highly risk-averse contractor is profitable with the optimal contract. The contractor is unprofitable in three sets of draws with the lump-sum-only contract. The contractor's expected profit facing the

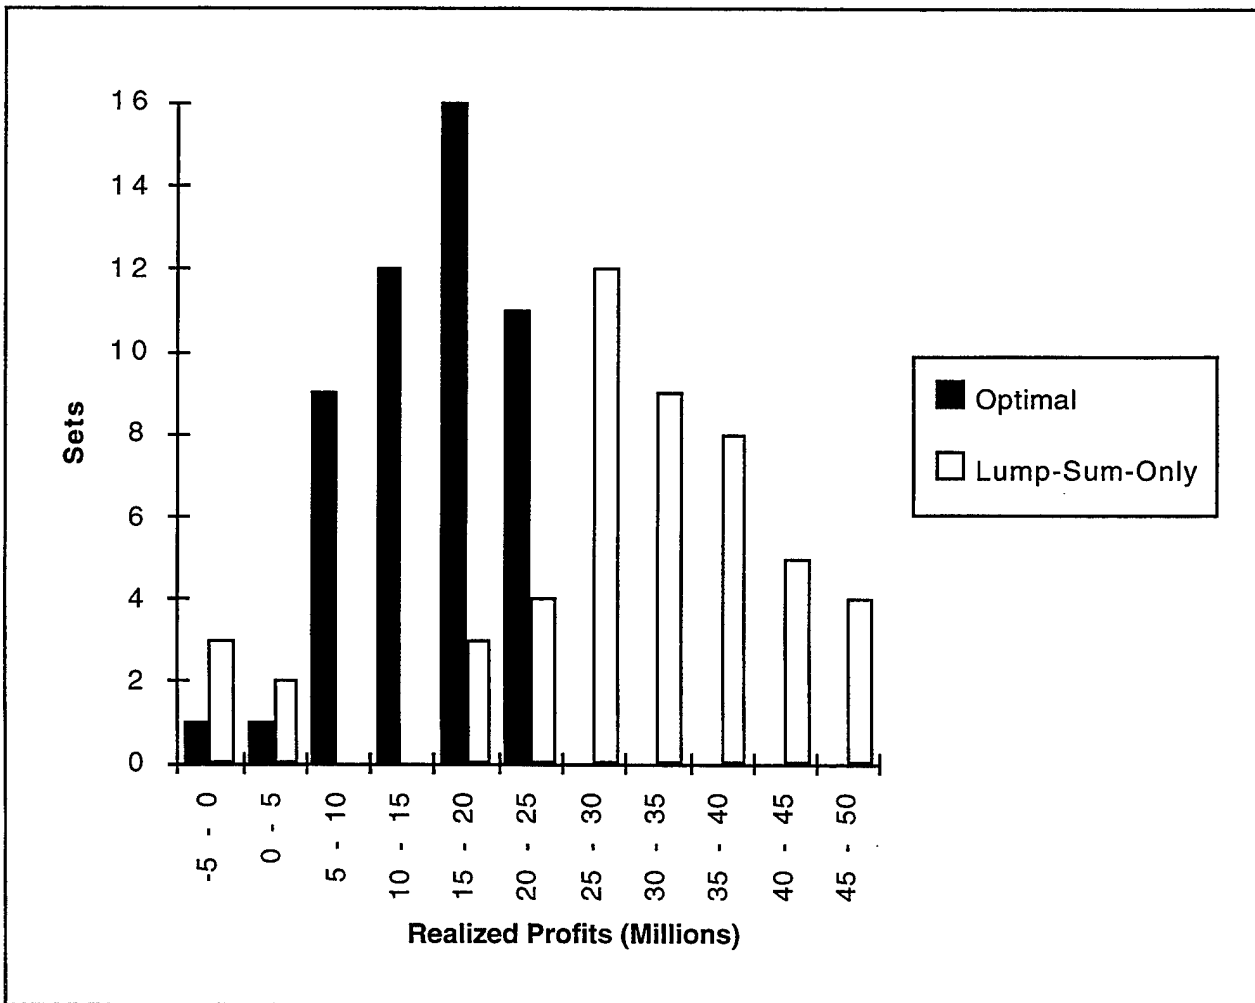


Figure 3—Risk-Averse Contractor Profits with Optimal and Lump-Sum-Only Contracts

Table 3 optimal contract is \$14.8 million. By contrast, the contractor's expected profit with the lump-sum-only contract is \$29.7 million. Of course, the contractor's expected utility with either contract, by definition, is zero. The lump-sum-only contract implies the risk-averse contractor is bearing more risk, so the contractor must be compensated for doing so.

In conclusion, we see optimal repair contracts with risk-averse contractors involve risk-sharing between the contractor and the government ($c_2 > 0$). Such risk-sharing reduces the government's expected costs. However, complete risk elimination, as in Table 2's time and materials contract, is highly undesirable as the contractor has strong temptation to cheat on such a contract.

NO PER-UNIT PENALTY?

One unusual aspect of Table 3's optimal contract is the per-unit penalty, $c_1 < 0$. Under this type of contract, the contractor not only has to fix any unit that breaks, but must pay a penalty on top of any repair costs.

This penalty is part of the optimal contract because 88 percent spares cost-sharing dulls the contractor's incentives to undertake high-quality repairs. A penalty every time a part breaks increases the contractor's incentive to build high quality into repairs.

Perhaps, however, a per-repair penalty is not practical. As an alternative, we considered a case where c_1 is constrained to be zero and the lump-sum and the spares cost-sharing, c_2 , can vary. Table 4 shows the optimal contract in this case.

As expected, the government's expected expenditure is greater in this case compared to Table 3's. However, the increase in the government's expected expenditure is less than 2 percent. Without a negative c_1 , spares cost-sharing, c_2 , falls. Table 4 involves 84 percent, rather than 88 percent, cost-sharing.

The broader message of Table 4, however, is that a per-repair penalty ($c_1 < 0$) does not appear to be centrally important. One can achieve comparable results with $c_1 = 0$ and somewhat less spares cost-sharing.

Table 4
Simulation Contract Approaches with Risk-Averse Contractor
98% Required Aircraft Availability
 $c_1 = 0$ Contract

	Parameter	$c_1 = 0$ Contract
Government	c_1	\$0
	c_2	\$166,786
	Lump-sum	\$67,734,400
Contractor	x	5.39
	q	1.00
	m	0.49
	t_r	1
Government	Expected expenditure	\$102,854,600

NOTE: Simulation based on 50 sets of 1250 periods. Expected contractor utility is zero under each contract.

For the rest of this report, we allow c_1 to be negative. Table 4 suggests, however, that this is not a critical feature of repair contracts.

THE COST-AVAILABILITY CURVE

As noted, Table 3 displays the optimal contract when 98 percent of the aircraft must be available in every period.

We can use the same simulation procedure to sketch the cost-availability frontier, i.e., a display of the costs of different aircraft availability levels. Table 5 shows the optimal contracts and the results of a lump-sum contract for various availability requirements.

Not surprisingly, the government's expected expenditures with the respective optimal contracts are increasing in the required aircraft availability rate. The cost per available aircraft also increases with the required availability rate. We see that 100 percent availability costs more than three times as much as 50 percent availability. The lump-sum-only approach is least appropriate at high availability levels. With a low enough availability requirement, the contractor's risks are not large, so the lump-sum-only approach is more appropriate.

Table 5
Optimal Contracts and Required Aircraft Availability Levels

Case	c1	c2	Lump-sum	x	m	Expected Expenditure
50% Optimal	\$7,877	-\$79,399	\$19,102,000	2.95	0.54	\$29,336,400
50% Lump-Sum-Only	\$0	\$0	\$31,318,300	2.67	0.54	\$31,318,300
60% Optimal	-\$2,000	\$170,591	\$42,115,600	3.26	0.50	\$40,584,300
60% Lump-Sum-Only	\$0	\$0	\$47,126,200	6.33	0.65	\$47,126,200
70% Optimal	-\$13,700	\$177,302	\$72,910,200	3.09	0.63	\$56,255,000
70% Lump-Sum-Only	\$0	\$0	\$66,629,300	7.89	0.67	\$66,629,300
80% Optimal	-\$7,577	\$180,599	\$69,555,000	3.82	0.58	\$68,483,800
80% Lump-Sum-Only	\$0	\$0	\$86,473,500	9.01	0.67	\$86,473,500
90% Optimal	-\$9,968	\$178,677	\$85,068,400	4.04	0.62	\$87,354,000
90% Lump-Sum-Only	\$0	\$0	\$106,840,100	10.16	0.67	\$106,840,100
95% Optimal	-\$11,972	\$178,945	\$94,893,200	4.09	0.64	\$97,468,300
95% Lump-Sum-Only	\$0	\$0	\$116,792,700	10.65	0.67	\$116,792,700
98% Optimal	-\$2,628	\$174,110	\$73,292,700	5.13	0.53	\$101,025,900
98% Lump-Sum-Only	\$0	\$0	\$122,728,000	11.06	0.67	\$122,728,000
100% Optimal	-\$7,516	\$181,010	\$87,069,100	4.74	0.59	\$103,901,700
100% Lump-Sum-Only	\$0	\$0	\$126,699,900	11.28	0.67	\$126,699,900

NOTE: Simulation based on 50 sets of 1250 periods. Expected contractor profits are zero under each contract. For each case, $q=1$ and $t_r=1$.

Figure 4 plots the cost-availability curve.

It surprised us that the cost-availability curves in Figure 4 are so linear. Our intuition from inventory theory was that the curves would be more convex, i.e., the cost of 100 percent aircraft availability would be extraordinary. The key to approximate linearity in Figure 4, it appears, is that repair capacity (x), as well as the number of spares purchased, is allowed to vary with the required availability level. Increasing the availability requirement, holding the capacity level fixed would result in a convex increase in costs as only spares could be used to obtain an increased availability level.

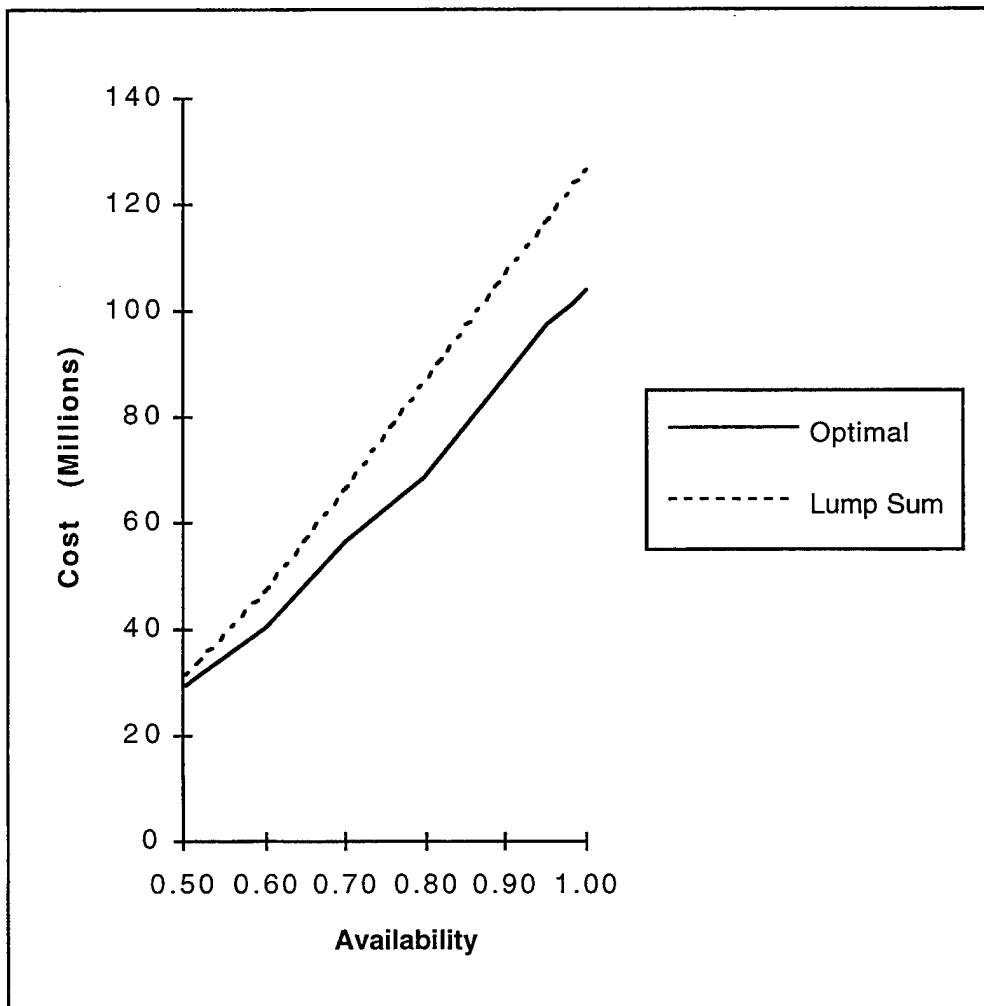


Figure 4—Cost-Availability Curve

With both the optimal contract and, even more markedly, the lump-sum-only contract, the risk-averse contractor chooses to have a great deal of excess repair capacity on average. For example, in the 100 percent availability case, the optimal contract results in an average repair capacity utilization rate of 42.8 percent. In the same case, the lump-sum-only contract results in an average repair capacity utilization rate of 14.9 percent. The contractor is quite averse to having to purchase expensive spares.

This excess capacity cushions the contractor against the variability in demands. One would see considerable convexity if one had to purchase spares to buy out the variability at high availability rates.

OPTIMALITY AND ROBUSTNESS

The "optimal" contract in Table 3 is only optimal, according to our simulation, if all the myriad parameters and assumptions underlying Table 3 hold true. Further, the Table 3 optimal contract is based on a finite number of random draws.

As a practical matter, we may not so much be interested in a specifically optimal contract like the Table 3 optimal contract. Instead, we may wish to find a contract that is effective under a variety of circumstances, especially when we do not know all the detailed parameters of the model exactly. We refer to such a versatile contract as being robust. A robust contract is optimal in that it minimizes expected costs, given government uncertainty about certain parameters that characterize the world.

Chapter 3 discusses our analysis of contract robustness. It turns out that Table 3's optimal contract is surprisingly robust. It closely resembles the optimum contract even with some major perturbations in the model's parameter values.

3. EFFECTS OF CHANGING ASSUMPTIONS

Table 3's findings apply to a very specific set of parameters and a specific type of contract. In this chapter, we examine how optimal contracts change as situations change. We also evaluate whether the Table 3 optimal contract performs adequately in these revised cases. Each case in this chapter uses the 98 percent required availability standard.

We start with a change to the contractor's objective function. Suppose the contractor is a risk-neutral profit-maximizer. How should the government design contracts for such a contractor?

Another likely possibility is that aircraft parts fail in a manner other than the failure distribution assumed in Table 3. We examine perturbations to Table 3's scenario whereby, for instance, parts fail more or less frequently. We assume the contractor knows the true failure rates, but the government may be misinformed. As the government might not anticipate failure rates correctly, it is important, too, to note how the Table 3 optimal contract works with different part failure distributions.

There are other possible perturbations to the Table 3 scenario that the government is more likely to know about. We examine cases where modifications are prohibitively expensive, where speeding hands-on repair time is costly, and where spares are less costly.

Finally, the government itself might wish to change the contracting approach. We examine a case of a one-year contract (our base case is a five-year contract) and a case of having ten times fewer aircraft repaired under a contract.

A RISK-NEUTRAL CONTRACTOR

Table 3 assumes a risk-averse contractor. Suppose, instead, the contractor were risk-neutral. Clearly, risk-neutral contractors may not be likely, but it is possible that a contractor would not be as risk-averse as the Table 3 contractor. Hence, it is valuable to examine ramifications of loosening the risk-aversion assumption. Under risk-

neutrality, the contractor's utility function has the form $U(k)=k$ for all values of k . A risk-neutral contractor weights a gain of $\$k$ as favorably as it weights a loss of $\$k$ unfavorably. Table 6 shows the optimal contract in this case.

Table 6 shows that a contract that relies solely on a lump-sum payment to the risk-neutral contractor is essentially optimal. In contrast, Table 3's optimal contract is problematic with a risk-neutral contractor. In our model, a risk-neutral contractor would choose to build too little capacity and would choose too low an m value while relying excessively on the subsidized spares to maintain the required aircraft availability level. If the government is wrong about the degree of contractor risk-aversion, fine-tuned contracts like the Table 3 contract can be troublesome. Of course, there is a large difference in the risk-aversion levels in Table 3 and Table 6, so perhaps it is comforting that Table 3's optimal contract is not even more inappropriate in this case.

Figure 5 shows the frequency of different ranges of realized contractor profits from 50 sets of simulations of Table 6's lump-sum-only contract with a risk-neutral contractor. Note Figure 5 is on a

Table 6
Simulation Contract Approaches with Risk-Neutral Contractor
98% Required Aircraft Availability

	Parameter	Optimal Contract	Table 3 Optimal Contract	Lump-Sum-Only Contract
Government	$c1$	\$588	-\$2,628	\$0
	$c2$	-\$3,224	\$174,110	\$0
	Lump-sum	\$82,078,900	\$57,362,200	\$82,868,900
Contractor	x	5.68	4.21	5.68
	q	1.00	1.00	1.00
	m	0.62	0.48	0.62
	t_r	1	1	1
Government	Expected expenditure	\$82,864,400	\$93,530,400	\$82,868,900

NOTE: Simulation based on 50 sets of 1250 periods. Expected contractor profits are zero under each contract.




Figure 5—Risk-Neutral Contractor Profits with Lump-Sum-Only Contract

different scale than Figure 3. In contrast to Figure 3, there is a substantial chance the risk-neutral contractor loses money. Of course, with contractor risk-neutrality, the government has the luxury of being able to set contractor expected profits to zero.

Comparing Tables 3 and 6, the government spends more money when facing a risk-averse, rather than a risk-neutral, contractor because the government must compensate the risk-averse contractor for bearing risk. Table 3's lump-sum contract carries an expected government expenditure of \$122,728,000. In contrast, Table 6's lump-sum contract has an expected government cost of \$82,868,900. Hence, if one uses lump-sum

contracts, Table 3's contractor risk-aversion costs the government \$39,859,100, or a 48 percent increase in expected costs. Table 3's optimal contract eliminates \$21,702,100 of this risk-aversion cost, or 54.4 percent of the expected cost of risk-aversion. Hence, Table 3's optimal contract is fairly effective in reducing the government's cost of risk-aversion, but a 22 percent risk-aversion premium remains.

GOVERNMENT HOLDS INCORRECT DISTRIBUTIONAL BELIEFS

The results in Table 3 are predicated on some heroic informational assumptions; e.g., the government and the contractor know part failure fractions are drawn from a lognormal distribution with mean 0.013 and variance twice as large as the mean, truncated at 1.

Suppose, for instance, that the risk-averse contractor knew the true statistical distribution of part failure fractions was lognormal with mean 0.0065 (half the failure fraction the government expects) and variance twice as large as the mean, truncated at 1. Table 7 illustrates the truly optimal contract in this case (which the government does not know) as well as the results if the government were to offer Table 3's optimal contract or the lump-sum-only contract.

Table 7

**Simulation Contract Approaches with Risk-Averse Contractor
Failure Fraction Half Expected Level
98% Required Aircraft Availability**

	Parameter	Optimal Contract	Table 3 Optimal Contract	Lump-Sum-Only Contract
Government	<i>c</i> 1	-\$2,567	-\$2,628	\$0
	<i>c</i> 2	\$177,777	\$174,110	\$0
	Lump-sum	\$50,617,500	\$52,301,400	\$111,564,300
Contractor	<i>x</i>	3.93	3.90	7.57
	<i>q</i>	1.00	1.00	1.00
	<i>m</i>	0.31	0.32	0.32
	<i>t_r</i>	1	1	1
Government	Expected expenditure	\$82,540,100	\$82,905,000	\$111,564,300

NOTE: Simulation based on 50 sets of 1250 periods. Expected contractor utility is zero under each contract.

Table 7 shows that the optimal contract with half the expected failure rate is very similar to the Table 3 optimal contract. As in Table 3, forcing all risk on the risk-averse contractor through a lump-sum-only contract is costly to the government. Note the risk-averse contractor chooses far more repair capacity when facing the lump-sum-only contract than the optimal contract. Due to this contractor risk-aversion, the government pays a 35 percent premium with a lump-sum-only contract in this case.

Table 8 presents the reverse case, i.e., a part failure fraction twice what the government expects. Again, the optimal contract in Table 8 closely resembles the Table 3 optimal contract, and the Table 3 optimal contract performs well.

Comparing Tables 3, 7, and 8, we see that changes in the expected part failure rates do not proportionally change the government's expected costs. Table 7's halving of the part failure rate reduces the government's expected expenditure less than 20 percent. Table 8's doubling of the part failure rate increases the government's expected expenditures only by roughly 25 percent. There are important fixed repair costs in this parameterization of the model. This result

Table 8
Simulation Contract Approaches with Risk-Averse Contractor
Failure Fraction Twice Expected Level
98% Required Aircraft Availability

	Parameter	Optimal Contract	Table 3 Optimal Contract	Lump-Sum-Only Contract
Government	<i>c</i> 1	-\$2,723	-\$2,628	\$0
	<i>c</i> 2	\$185,535	\$174,110	\$0
	Lump-sum	\$103,291,000	\$105,446,000	\$141,472,000
Contractor	<i>x</i>	9.07	9.09	9.31
	<i>q</i>	1.00	1.00	1.00
	<i>m</i>	0.67	0.67	0.71
	<i>t_r</i>	1	1	1
Government	Expected expenditure	\$126,305,400	\$126,854,600	\$141,472,000

NOTE: Simulation based on 50 sets of 1250 periods. Expected contractor utility is zero under each contract.

suggests that if the government wants to tie payments to the contractor to the number of flying hours, for instance, the payment schedule should be nonlinear. The contractor should receive a fixed payment along with a payment that varies with the number of flying hours.

PROHIBITIVELY EXPENSIVE MODIFICATIONS

The cases presented heretofore are characterized by $m > 0$ in equilibrium; i.e., parts that break are repaired to be "better than new." Perhaps such thorough repairs are not feasible.

Looking at the other extreme, we examined the case where both the fixed and marginal costs of modifications were 100 times larger; i.e., it was prohibitively expensive even to consider modifications. Table 9 shows the resultant optimal contract.

Table 9's optimal contract closely resembles Table 3's. Again, we see a combination of a per-repair penalty ($c1 < 0$) and sizable spares cost-sharing is optimal. Despite the major limitation in repair technology, the Table 3 optimal contract continues to work well in this case. As in Table 3, the lump-sum-only contract induces the risk-averse contractor to build excess repair capacity.

Table 9

**Simulation Contract Approaches with Risk-Averse Contractor
Prohibitively Expensive Modifications
98% Required Aircraft Availability**

	Parameter	Optimal Contract	Table 3 Optimal Contract	Lump-Sum-Only Contract
Government	$c1$	-\$3,132	-\$2,628	\$0
	$c2$	\$180,616	\$174,110	\$0
	Lump-sum	\$82,638,400	\$84,728,400	\$203,090,700
Contractor	x	5.46	5.42	10.39
	q	1.00	1.00	0.83
	m	0.00	0.00	0.00
	t_r	1	1	1
Government	Expected expenditure	\$153,250,600	\$155,198,200	\$203,090,700

NOTE: Simulation based on 50 sets of 1250 periods. Expected contractor utility is zero under each contract.

When modifications are no longer an option, the cost of the optimal contract is over 1.5 times Table 3's optimal contract cost; more parts need to be repaired. Further, the contractor's capacity level needs to increase ($x=5.46$ versus $x=5.13$ in Table 3).

The repair stand utilization rate with Table 9's optimal contract is approximately 76 percent. Since $m=0$, the expected repair rate is constant over time. However, given the stochastic nature of part failure, it proves to be optimal to have some unused capacity, on average, because having to use spares to maintain availability is very expensive.

HANDS-ON REPAIR SPEED IS COSTLY

Our results heretofore assume the contractor can choose any repair speed ("hands-on time") between 1 and 50 days with no inherent cost associated with this choice.

An opposing view is that it does cost a contractor to reduce hands-on repair time; e.g., capital expenditures may be required. Hence, we investigated how the optimal contract and the associated outcome changes over a couple of hands-on repair speed cost scenarios.

Table 10 outlines the case where there is an incremental cost to the contractor of \$500,000 per day for hands-on repair time faster than 50 days.

In Table 10's \$500,000 per day case, very fast ($t_r=1$) hands-on repair time is chosen in the optimal contract, even with a \$24,500,000 penalty to the contractor for choosing this speed level. The Table 3 contract remains virtually optimal.

Table 11, however, shows that eventually the costs of repair speed grow prohibitive. In Table 11's case, hands-on repair speed faster than 50 days costs \$1,500,000 per day. We now see an equilibrium with very slow repair. Further, repair capacity and the modification level fall (from $x=5.01/m=0.54$ in Table 10 to $x=4.23/m=0.30$ in Table 11), so spares are proportionally more important in maintaining the required aircraft availability level. As above, however, the Table 3 contract is essentially optimal in this case.

Table 10

Simulation Contract Approaches with Risk-Averse Contractor
 Hands-On Repair Speed Costs \$500,000 per Day
 98% Required Aircraft Availability

		Table 3		
Parameter		Optimal Contract	Optimal Contract	Lump-Sum-Only Contract
Government	<i>c1</i>	-\$3,608	-\$2,628	\$0
	<i>c2</i>	\$176,418	\$174,110	\$0
	Lump-sum	\$99,911,600	\$97,704,200	\$147,251,800
Contractor	<i>x</i>	5.01	5.13	10.99
	<i>q</i>	1.00	1.00	1.00
	<i>m</i>	0.54	0.53	0.67
	<i>t_r</i>	1	1	1
Government	Expected expenditure	\$125,153,500	\$125,412,100	\$147,251,800

NOTE: Simulation based on 50 sets of 1250 periods. Expected contractor utility is zero under each contract.

Table 11

Simulation Contract Approaches with Risk-Averse Contractor
 Hands-On Repair Speed Costs \$1,500,000 per Day
 98% Required Aircraft Availability

		Table 3		
Parameter		Optimal Contract	Optimal Contract	Lump-Sum-Only Contract
Government	<i>c1</i>	-\$3,183	-\$2,628	\$0
	<i>c2</i>	\$178,747	\$174,110	\$0
	Lump-sum	\$78,265,900	\$79,917,000	\$193,708,000
Contractor	<i>x</i>	4.27	4.31	4.10
	<i>q</i>	1.00	1.00	1.00
	<i>m</i>	0.30	0.30	0.59
	<i>t_r</i>	50	50	50
Government	Expected expenditure	\$155,636,300	\$156,459,600	\$193,708,000

NOTE: Simulation based on 50 sets of 1250 periods. Expected contractor utility is zero under each contract.

It turns out that with the parameters we have used, the crossover point from a $t_r=1$ equilibrium to a $t_r=50$ equilibrium occurs at a cost per day of slightly over \$1,000,000.

SPARES HALF AS COSTLY

We also investigated a case where spares were half as costly as before (\$99,289 rather than \$198,578), but everything else (e.g., repair costs, failure rates) was held constant. Table 12 presents the optimal contract in this case.

Table 12's optimal contract has the magnitude of c_2 , the spares cost-sharing, reduced somewhat more than proportionally to the reduction in the spares cost. Table 3's optimal contract had 88 percent spares cost-sharing; Table 12's has 77 percent spares cost-sharing. Table 12's optimal contract's cost is 80 percent as large as Table 3's. Not surprisingly, we see somewhat less reliance on repair capacity ($x=5.05$ versus $x=5.13$ in Table 3) with the halving of the relative price of spares.

It is also not surprising that Table 3's optimal contract does not perform well in this situation. The contractor eagerly lives off

Table 12
Simulation Contract Approaches with Risk-Averse Contractor
Spares Cost Half as Much
98% Required Aircraft Availability

	Parameter	Optimal Contract	Table 3 Optimal Contract	Lump-Sum-Only Contract
Government	c_1	-\$3,305	-\$2,628	\$0
	c_2	\$76,491	\$174,110	\$0
	Lump-sum	\$74,718,900	-\$285,675,000	\$93,599,800
Contractor	x	5.05	0.00	5.35
	q	1.00	NA ¹	1.00
	m	0.54	0.00	0.57
	t_r	1	NA	1
Government	Expected expenditure	\$81,242,800	\$578,819,900	\$93,599,800

NOTE: Simulation based on 50 sets of 1250 periods. Expected contractor utility is zero under each contract.

¹The quality and speed of repair are irrelevant if $x=0$. Meanwhile, one does find $m=0$ in this situation since there are fixed costs associated with any positive modification level.

\$99,289 spares for which he receives \$174,110 each. In prior cases, we have seen that informational asymmetries, e.g., about failure rates or modification costs, are not always costly to the government. The Table 3 optimal contract continues to work fairly well. Table 12 presents the (perhaps obvious) caveat that if the government chooses to mitigate contractor risk by subsidizing expensive spare parts, the government does have to have a reasonable understanding of the costs of these spares.

DIFFERENT CONTRACTING APPROACHES

One-Year Contract

Table 3's contract runs five years. Suppose, instead, the government wished to have a one-year repair contract. Given annual budget cycles, the government often issues one-year contracts. Even government multiyear contracts generally involve annual renewal options unilaterally held by the government. Table 13 shows the optimal one-year contract.

The optimal contract in this case involves less risk-sharing than in Table 3 (81 percent down from 88 percent in Table 3). Table 3's optimal contract induces too little repair capacity and too small a

Table 13

**Simulation Contract Approaches with Risk-Averse Contractor
One-Year Contract
98% Required Aircraft Availability**

	Parameter	Optimal Contract	Table 3 Optimal Contract	Lump-Sum-Only Contract
Government	<i>c</i> 1	-\$2,571	-\$2,628	\$0
	<i>c</i> 2	\$160,250	\$174,110	\$0
	Lump-sum	\$36,227,900	\$29,183,800	\$86,421,400
Contractor	<i>x</i>	18.85	3.88	19.08
	<i>q</i>	1.00	1.00	1.00
	<i>m</i>	0.31	0.00	0.20
	<i>t</i> _r	1	1	1
Government	Expected expenditure	\$54,772,000	\$77,685,000	\$86,421,400

NOTE: Simulation based on 50 sets of 250 periods. Expected contractor utility is zero under each contract.

modification level. Consequently, distended government costs occur due to excessive spares reliance.

The government's expected cost of the optimal one-year contract is more than 50 percent of Table 3's five-year expected costs. The short duration of the contract implies a low m ; consequently, the cost per year is considerably greater.

The average capacity utilization rate in the optimal case is about 18 percent. The considerable variance in part failure requires large excess repair capacity, on average.

Contractor Responsible for Ten Times Fewer Aircraft

The government might want to have fewer aircraft repaired by the contractor. Hence, we investigated the case of ten times fewer aircraft, i.e., 40.3 aircraft of which 39.5 must be available. All other parameters were kept the same, most notably the fixed costs of a given modification level. Table 14 shows the resultant optimal contract.

Table 14 shows that aircraft modification is no longer desirable with fewer aircraft. However, the major change in the number of

Table 14

**Simulation Contract Approaches with Risk-Averse Contractor
Contractor Responsible for Ten Times Fewer Aircraft
98% Required Aircraft Availability**

		Table 3		
Parameter		Optimal Contract	Optimal Contract	Lump-Sum-Only Contract
Government	$c1$	-\$2,804	-\$2,628	\$0
	$c2$	\$182,412	\$174,110	\$0
	Lump-sum	\$7,865,800	\$8,371,000	\$17,345,700
Contractor	x	0.55	0.54	0.83
	q	1.00	1.00	1.00
	m	0.00	0.00	0.60
	t_r	1	1	1
Government	Expected expenditure	\$14,589,600	\$14,866,400	\$17,345,700

NOTE: Simulation based on 50 sets of 1250 periods. Expected contractor utility is zero under each contract.

aircraft notwithstanding, the Table 3 contract structure continues to work well, as long as the lump-sum payment shrinks appropriately.

Table 14 suggests considerable economies of scale in aircraft repair, given the cost function we have assumed. Specifically, modifications are no longer cost-effective at this smaller scale. Hence, despite having 10 percent as many aircraft, Table 14's minimized expected government expenditure is over 14 percent as large as Table 3's.

4. CONCLUSION AND IMPLICATIONS

This research has developed and simulated a model of repair contractor behavior. Models are, by definition, oversimplifications of the real world. Hence, we urge against interpreting these results too literally. However, we believe there are a number of tentative lessons from these models and implementations. First, per-repair contracts often seem to provide poor incentives. The contractor has incentive to choose slow, low-quality repair. Hence, the repair contractor has to be intensively monitored to provide reasonable repair speed and quality.

According to our simulation, the optimal contract combines a lump-sum payment, expensive item cost-sharing, and an availability guarantee. With this sort of contract, the contractor guarantees a specified availability level. This type of arrangement appears to provide excellent incentives to contractors. The Air Force is using an arrangement of this sort for C-21 support, while the Navy has a contract of this sort for LN-15C inertial navigation unit repair.

Across a fairly wide variety of scenarios, such contracts will give contractors incentives to provide fast, high-quality repair without explicit monitoring of or rewards for contractor capacity, quality, or speed. Instead, the government only needs to make sure aircraft availability guarantees are being fulfilled. With thoughtfully constructed contracts, the government's monitoring costs should be comparatively minimal. Intensive auditing and monitoring of contractors may not be an inevitable component of government contracting.

Such contracts do, however, require a well-informed contractor. Hence, contracts with lump-sum payments, expensive item cost-sharing, and availability guarantees seem most applicable for mature weapon systems with predictable failure patterns.

APPENDIX A: BASIC MODEL AND ENHANCEMENTS

BASIC MODEL

The model accommodates a single type of part. It operates under the assumption that the buyer will operate N of these parts in every period. For a new part, there is a stochastic failure fraction ε_k in period k . ε_k is a random variable drawn from some distribution $g(\varepsilon)$, where $g(\varepsilon)$ has support over the interval $[0,1]$.

For period 1, $\varepsilon_1 N$ parts are broken and go to the contractor for repair. In this model, there is one repair contract, though multiple contractors may compete for this contract. There is no organic repair in this model.

The repair contractor must choose some repair capacity x before the contract begins. If $\varepsilon_1 N < x$, the parts will be fixed for next period. If not, new parts will be needed to get the number of operating parts back up to N for next period. We assume N parts are operated every period. The contractor must purchase extra parts to the extent there are not enough parts repaired. The contractor may, however, be reimbursed, at least partially, for these purchases.

This model is based on a fixed-availability assumption. The contractor promises to provide a fixed level of aircraft availability to the government. This is not simply a model in which the contractor tries its best to maximize availability. Instead, the contractor promises to purchase spares as needed to maintain the required availability level.¹

The contractor also must choose the quality, q , of its repairs, $q \in [0,1]$. If the contractor chooses quality level q , the newly repaired

¹One might justifiably argue that it is unrealistic that a contractor could instantaneously acquire needed spares. A related model would have the contractor pay a penalty equal to the cost of a spare every time a spare would be needed to hit the availability target. Such a penalty-based model would differ from this chapter's model, though, in that actual aircraft availability would not always hit the target. Also, the government's net expenditure would be lower as it would receive penalty revenue.

part fails during period 2 with probability $(1-q+q\epsilon_2)$. Any value of q can be chosen between 0 and 1, inclusive. If the contractor chooses $q=0$, the part will certainly fail next period. If the contractor chooses $q=1$, the part will fail next period with probability ϵ_2 , the failure probability of a new part. If the repaired part does not fail in its first period after repair, its failure probability thereafter corresponds to the failure probability of a new part.

Given this model, we can define five state variables and their interrelationships. Define $F(k)$ to be the number of parts working in period k immediately out of repair, i.e., the number of parts that just got fixed. Define $P(k)$ to be the number of new parts brought into service in period k (due to an inadequate number of parts coming out of repair). Define $L(k)$ to be the number of leftover parts in period k . With exactly N parts operating every period, it is possible that there will be unused, but operable, parts. Define $B(k)$ to be the number of parts newly broken as a result of operation in period k . Finally, define $D(k)$ to be the number of parts in the contractor's depot or repair system after period k . $D(k)$ will exceed $B(k)$ if the contractor has unfixed, broken parts overhanging from prior periods.

Expressions for these state variables and the explanations for their derivation are found below.

$$\begin{aligned}F(k) &= \min(D(k-1), x) \\P(k) &= \max(B(k-1) - F(k) - L(k-1), 0) \\L(k) &= \max(F(k) + L(k-1) - B(k-1), 0) \\B(k) &= (N - F(k))\epsilon_k + F(k)(1 - q + q\epsilon_k) \\D(k) &= \max(D(k-1) - x, 0) + B(k).\end{aligned}$$

First, $F(k)$ is derived by noting that, by definition, $D(k-1)$ parts were in the repair system during period $k-1$. If $D(k-1) < x$, where x is the repair capacity, then all $D(k-1)$ parts are fixed. If not, only x are fixed.

$P(k)$ is the number of new parts brought into service when the repair system can't keep up with breakage. The Air Force loses $B(k-1)$ parts that broke in period $k-1$. The Air Force gains $F(k)$ parts, which

emerge from the repair system and can draw on any other unused parts. Hence, if $B(k-1) > F(k) + L(k-1)$, the contractor needs to purchase $B(k-1) - F(k) - L(k-1)$ new parts to cover the shortfall in period k .

$L(k)$ is the number of leftover, unused parts in period k . If the number of newly fixed parts, $F(k)$, plus leftover parts, $L(k-1)$, exceeds the previous period's breakage, $B(k-1)$, one will have leftover parts.

$B(k)$ is the number of broken parts in period k . Newly repaired parts, $F(k)$, break with probability $(1 - q + q\epsilon_k)$. Any parts that were just fixed fly this period, we assume. Other parts that fly, $N - F(k)$, break with probability ϵ_k .

$D(k)$ is the number of parts in the repair system after period k . That includes parts broken in period k ($B(k)$), plus any leftover parts if the system capacity, x , wasn't sufficient to fix all of period $k-1$'s parts in the repair system, $D(k-1)$.

Figure 6 sketches this model.

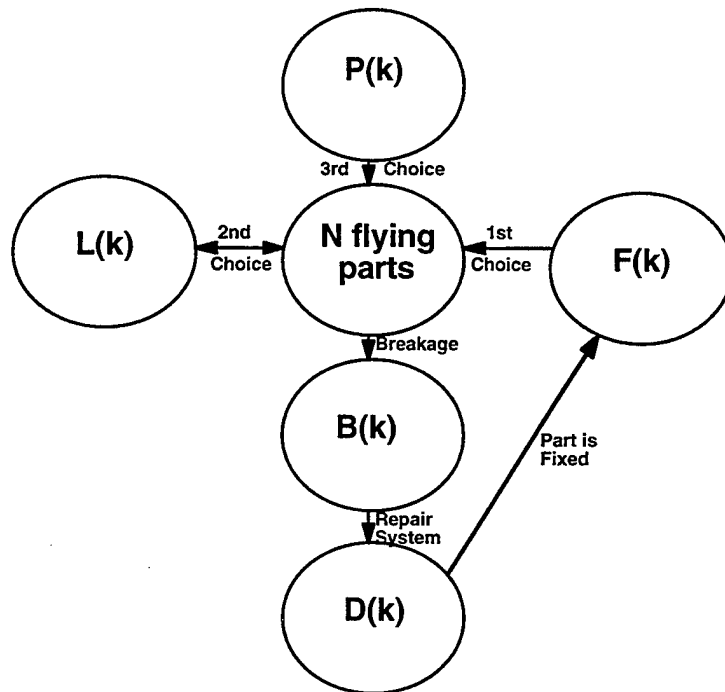


Figure 6-Basic Model

The contractor chooses capacity, x , and quality, q . The contractor is assumed to have costs that are a function of $F(k)$, x , q , and $P(k)$, if the contractor both fixes parts and builds new ones. (If the contractor only does repairs, $P(k)$ is not in the contractor's cost function.) Meanwhile, the contractor receives payment from the government based on $F(k)$ and $P(k)$. For any period k , the contractor's marginal revenue is

$$c_1 * F(k) + c_2 * P(k),$$

while the contractor's marginal cost is

$$(\text{Cost of a repair, a function of } q) * F(k) + (\text{Cost of a spare}) * P(k).$$

The contractor also has fixed costs (for capacity, x) and may receive a lump-sum payment. Given these revenues and costs, the contractor will choose x and q to maximize its expected utility.

Meanwhile, the government will want to offer a contract that minimizes its expected expenditures on repairs and new parts, taking into account the contractor maximizing its utility in its choices of x and q .

MODEL ENHANCEMENTS

There are a variety of realism enhancements one can build into this model. One important issue is whether the repair contractor will choose to do modifications to parts in the repair process that make them "better than new"; i.e., modified parts have failure probabilities lower than new parts.

To incorporate modifications into this model, suppose the contractor chooses the modification level m , $m \in [0,1]$. After a repair, the part's failure probability henceforth is $(1-m)$ times the old failure probability. If $m=0$, there is no change in the part's failure probability. If $m=1$, the part will never fail again.

With modification in the model, there is a considerably larger number of state variables. There are "old," i.e., unmodified, parts as well as modified parts. Conceptually, however, the model is the same.

Hands-on repair time can also be a contractor choice. (In the basic model, all repairs take one period to accomplish, if capacity is

available.) Let us suppose repairing an already modified, but broken, part takes t_r periods while modifying a part takes t_m periods. (If $m=0$, no modification occurs. In this case, t_m is the time it takes to repair a part the first time it is in the repair system and t_r is the time subsequent repairs take. One way to handle this would be to have $t_m=t_r$ for sufficiently small values of m .) t_r and t_m could be exogenous or they could be contractor choice variables. The parameters t_r and t_m do not include any time spent in repair queues. Clearly, queue time will be a function of t_r , t_m , x , q , m , and the aircraft failure distribution.

With the hands-on portion of repairs and modifications taking more than one period, it becomes an issue how the capacity constraint binds. We assume only x units can begin the repair process in a given period.²

We might also suspect that repair facilities themselves might not always work. Suppose there is a stochastic failure fraction λ_k for repair capacity in period k , where λ_k is a random variable drawn from some distribution $h(\lambda)$ with support over the interval $[0,1]$. λ_k could be correlated with ϵ_k . Effective repair capacity in period k is $x(1-\lambda_k)$. We assume the contractor has no control over λ_k .

To incorporate modifications, repair lags, and repair facility failures, define $OP(k)$ to be the number of "old" parts operating in period k . Define $MP(k)$ to be the number of modified parts operating in period k . If there are extra parts, one always operates all available modified parts before operating any old parts. Define $OB(k)$ to be the number of old parts broken during operation in period k . Define $MB(k)$ to be the number of modified parts broken during operation in period k . Define $OD(i,k)$ to be the number of old parts in the i th stage of repair in period k , $i \leq t_m$. (If there is not enough repair capacity to start repair on a part, it stays in the first stage of repair until capacity becomes available.) Define $MD(i,k)$ to be the number of modified, but broken, parts in the i th stage of repair in

²Below, we outline the case where x is the maximum number of units at any stage of the repair process at any point in time. The simulation results, however, are based on having x be the maximum number of units that can start repair in a period.

period k , $i \leq t_r$. We assume the contractor will first fix broken parts that have already been modified. Define $F(k)$ to be the number of newly fixed parts in period k . Parts that are fixed are modified, but it is possible that a part will be fixed that had previously been modified. Define $P(k)$ to be the number of new, but not modified, parts the contractor needs to purchase in period k to assure that N parts operate in period k . Define $OL(k)$ to be the number of leftover (operable but not operating) old parts and $ML(k)$ to be the number of leftover modified parts. Finally, define $MT(k)$ to be the total number of modified, operable parts and $NM(k)$ to be the number of parts newly modified in period k .

Given these variables, we obtain the system

$$\begin{aligned}
 OP(k) &= N - MP(k) \\
 MP(k) &= \min(N, MP(k-1) - MB(k-1) + F(k)) \\
 OB(k) &= \varepsilon_k OP(k) \\
 MB(k) &= (1 - q + q(1 - m)\varepsilon_k)F(k) + (1 - m)\varepsilon_k(MP(k) - F(k)) \\
 OD(1, k) &= \max(OD(1, k-1) - \max(x(1 - \lambda_{k-1}) - MD(1, k-1), 0), 0) + OB(k) \\
 OD(2, k) &= \min(OD(1, k-1), \max(x(1 - \lambda_{k-1}) - MD(1, k-1), 0)) \\
 OD(i, k) &= OD(i-1, k-1), i \in [3, t_m] \\
 MD(1, k) &= \max(MD(1, k-1) - x(1 - \lambda_{k-1}), 0) + MB(k) \\
 MD(2, k) &= \min(MD(1, k-1), x(1 - \lambda_{k-1})) \\
 MD(i, k) &= MD(i-1, k-1), i \in [3, t_r] \\
 F(k) &= OD(t_m, k-1) + MD(t_r, k-1) \\
 P(k) &= \max(OB(k-1) + MB(k-1) - F(k) - OL(k-1) - ML(k-1), 0) \\
 OL(k) &= \max(OL(k-1) + OP(k-1) - OP(k) - OB(k-1), 0) \\
 ML(k) &= \max(MT(k) - MP(k), 0) \\
 MT(k) &= MT(k-1) + F(k) - MB(k) \\
 NM(k) &= OD(t_m, k-1).
 \end{aligned}$$

$OP(k)$ is the number of old parts operating in period k . We assume that, if available, one uses modified parts, $MP(k)$, rather than old parts. N parts are in use every period, so $OP(k) = N - MP(k)$.

$MP(k)$ is the number of operating, modified parts. N parts operate each period, so $MP(k)$ can't exceed N . $MP(k-1)$ modified parts operate in period $k-1$. $MB(k-1)$ of these break, but then $F(k)$ modified parts emerge from the repair system.

$OB(k)$ is the number of old parts broken in period k . $OP(k)$ old parts are in use in period k ; they break with probability ϵ_k .

$MB(k)$ is the number of modified parts broken in period k . $MP(k)$ modified parts operate in period k . $F(k)$ of these modified parts are newly fixed and modified; they fail with probability $(1-q+q(1-m)\epsilon_k)$. The rest, $MP(k)-F(k)$, fail with probability $(1-m)\epsilon_k$.

$OD(1,k)$ is the number of old parts in the first stage of the repair process in period k . In the first stage, no repair actually occurs unless repair capacity is available. Hence, parts can linger in the first stage. During period $k-1$, $OD(1,k-1)$ old parts are in the first stage of the repair process along with $MD(1,k-1)$ modified parts. We assume a repair algorithm in which you first fix already modified parts then fix old parts only if you have capacity. Hence, if $x(1-\lambda_{k-1}) > MD(1,k-1)$, some of the $OD(1,k-1)$ parts get fixed. If not, all $OD(1,k-1)$ are leftover and stay in the first stage. In addition, $OB(k)$ newly broken parts arrive.

$OD(2,k)$ is the number of old parts in the second stage of the repair process in period k . If $x(1-\lambda_{k-1}) > MD(1,k-1)$, there was some capacity available to old parts last period. $OD(1,k-1)$ old parts were available in the first stage last period. To the extent capacity existed, these parts become $OD(2,k)$.

$OD(i,k)$ is the number of old parts in the i th stage of the repair process in period k , $i \in [3, t_m]$. For these i values, the number in the i th stage of repair this period is simply the number in the $i-1$ st stage last period.

$MD(1,k)$ is the number of modified, but broken, parts in the first stage of the repair process in period k . If $MD(1,k-1) > x(1-\lambda_{k-1})$, there are unfixed, modified parts from previous periods. If not, there are only newly broken, modified parts.

$MD(2,k)$ is the number of modified, but broken, parts in the second stage of the repair process in period k . Last period, there were

$MD(1,k-1)$ modified, but broken, parts in the first stage. At most, $x(1-\lambda_{k-1})$ could actually enter the repair process and therefore hit the second stage.

$MD(i,k)$ is the number of modified, broken parts in the i th stage of the repair process in period k , $i \in [3, t_r]$. For these i values, the number in the i th stage of repair this period is simply the number in the $i-1$ st stage last period.

$F(k)$, as above, is the number of newly fixed parts in period k . $OD(t_m, k-1)$ formerly old, broken parts emerge from the repair system along with $MD(t_r, k-1)$ modified, formerly broken parts.

$P(k)$ is the number of new, but not modified, parts the contractor must purchase to insure that N parts operate in period k . $OB(k-1)$ old parts break and $MB(k-1)$ modified parts break. However, one has newly fixed parts, leftover old parts, and leftover modified parts to offset these breakages before purchases need to be made.

$OL(k)$ is the number of leftover old parts in period k . One might have leftover old parts from period $k-1$. One also has $OP(k-1)$ operating old parts in period $k-1$. $OB(k-1)$ of these break. $OP(k)$ old parts operate in period k . If $OP(k) + OB(k-1) < OL(k-1) + OP(k-1)$, one ends up with leftover old parts in period k .

$ML(k)$ is the number of leftover, modified parts in period k . We know $MT(k)$ parts are modified and operable while $MP(k)$ actually operate. The remainder are leftover.

$MT(k)$ is the total number of modified, operable parts. $MT(k-1)$ were around last period, but $MB(k)$ broke. $F(k)$ newly fixed parts come on-line, however.

Finally, $NM(k)$ is the number of newly modified parts that we know to be $OD(t_m, k-1)$.

Clearly, adding these enhancements to the model complicates the analytics of this modeling approach, but there is no major conceptual complication. As in the simpler case, the contractor would choose x , q , and now m , t_m , and t_r to maximize expected utility, and the government, in turn, would attempt to design a contract that minimizes its expected expenditures, given contractor behavior.

ANOTHER VIEW OF CAPACITY

As noted, in the current configuration of the model, x , capacity, is the maximum number of units that can start to be repaired in a period.

Another view of capacity is that x is the maximum number of parts that can be at any stage in the repair process at a point in time.

In such a case, the effective capacity facing parts awaiting repair

is $x(1 - \lambda_k) - \sum_{i=2}^{i_r} MD(i, k) - \sum_{i=2}^{i_m} OD(i, k)$. Hence, we would get

$$OD(1, k) = \max(OD(1, k-1) - \max(x(1 - \lambda_{k-1}) - \sum_{i=1}^{i_r} MD(i, k-1) - \sum_{i=2}^{i_m} OD(i, k-1), 0), 0) + OB(k)$$

$$OD(2, k) = \min(OD(1, k-1), \max(x(1 - \lambda_{k-1}) - \sum_{i=1}^{i_r} MD(i, k-1) - \sum_{i=2}^{i_m} OD(i, k-1), 0))$$

$$MD(1, k) = \max(MD(1, k-1) - x(1 - \lambda_{k-1}) + \sum_{i=2}^{i_r} MD(i, k-1) + \sum_{i=2}^{i_m} OD(i, k-1), 0) + MB(k)$$

$$MD(2, k) = \min(MD(1, k-1), x(1 - \lambda_{k-1}) - \sum_{i=2}^{i_r} MD(i, k-1) - \sum_{i=2}^{i_m} OD(i, k-1)).$$

B. THE CONTRACT SIMULATION PROGRAM

We developed a UNIX-based C computer program we called CONTRACT to undertake the simulations described in this paper. This appendix sketches the CONTRACT program.

The outermost loop of the program involves the government's search for the optimal contract. The government searches over c_1 , c_2 , and the lump-sum using the AMOEBA search algorithm described by Press et al. (1989) on pages 326-330. c_1 and c_2 can be negative or positive. The lump-sum payment is chosen to set the contractor's expected utility to zero. The government attempts to minimize its expected expenditures.

The AMOEBA search algorithm is not particularly efficient, but we chose it as a conservative algorithm that can hopefully find optima in very nonlinear situations. We were concerned that faster, derivative-based search algorithms, e.g., Davidon-Fletcher-Powell, would jump too quickly to solutions without surveying the set of options thoroughly.

Given a specification of the government's contract, the next loop of the CONTRACT program involves the contractor's search for its optimal parameters, e.g., capacity (x), quality (q), the modification level (m), and hands-on repair time (t_r). For each contractor capacity, quality, modification, and speed combination, the program computes the contractor's expected utility and the government's expected expenditure.

We used the AMOEBA search algorithm to find the contractor's optimal capacity, quality, and modification level, given the government's contract. Meanwhile, we discovered that repair time (t_r) always goes to a corner solution. Hence, we simply checked the corners of $t_r=1$ and $t_r=50$. (Obviously, in this loop, the contractor does not account for the fact that in the higher loop, the government adjusts the lump-sum payment to set the contractor's expected utility to zero. The lump-sum payment is exogenous from the contractor's perspective.)

We know a priori that $x \geq 0$, $0 \leq q \leq 1$, and $0 \leq m \leq 1$. In order to assure that these constraints held, the search algorithm actually searched for y_1 , y_2 , and y_3 , where $x = y_1^2$, $q = y_2^2 / (y_2^2 + 1)$, and $m = y_3^2 / (y_3^2 + 1)$.

The expected utility and expenditure values are computed from 50 sets of 1250 period simulations. For each period, there is a random draw that determines the part failure fraction (ϵ_k) and a different random draw that determines the repair station failure fraction (λ_k). The draws ϵ_k and λ_k come from a random number generator based on work by Marsaglia and Zaman (1991) and Feldman (forthcoming).

C. APPLICATION PARAMETERS

This appendix discusses our base-case parameters. These parameters are presented in Table 1 and used in the simulation results in Tables 2 and 3. These parameters are loosely based on the Coronet Deuce exercise described in Abell and Shulman (1992). However, this example is designed to be illustrative, not a tight representation of reality.

We assume a Radio Frequency (RF) repair stand costs \$5.35 million and lasts ten years. Hence, its annual fixed cost is \$535,000.¹ Repair stands are assumed to have a stochastic failure fraction drawn from a lognormal distribution with mean 0.1 and variance 0.2, truncated at 1.

Three F-16 parts are relevant to our analysis: the dual mode transmitter, the modular low power radio frequency (MLPRF), and the antenna. To accommodate the needs of the model, we will treat these as a single composite part with a pattern of failure and costs consistent with the underlying characteristics of these three parts. We assume that on a given flight hour, the dual mode transmitter fails with probability 0.005, the MLPRF fails with probability 0.006, and the antenna fails with probability 0.002. If failure probabilities are independent, the probability that nothing fails is 0.987. Hence, we assume part failure fractions per flight hour are drawn from a distribution with mean 0.013. We further assume this failure fraction distribution is lognormal and has variance twice as large as the mean, truncated at 1. We assume part failure and repair stand failure are statistically independent. These distributional assumptions describe the pre-modification failure pattern. If the contractor chooses to modify aircraft, the failure fraction falls.

¹For simplicity, we set aside problems associated with willingness to invest in a specific asset like an F-16 Radio Frequency repair stand or getting access to the technical data that allow a contractor to use such a specific asset. A more complete analysis would address the implications of these problems for contract design. For simplicity in this illustrative analysis, we assume that such an asset is freely available for rental at the annual fee noted.

We have data indicating that a dual mode transmitter costs \$199,778 new, a MLPRF costs \$229,134 new, and an antenna costs \$98,756.² Hence, the failure weighted average price of a new part is \$198,578. We will simplify our problem by assuming there is one "typical" part with this failure weighted average price.

We also have data indicating that dual mode transmitter repair has a marginal cost of \$6469.38, MLPRF repair has a marginal cost of \$5544.75, and antenna repair has a marginal cost of \$3265.59. Hence, the average marginal repair cost, conditional on some failure occurring, is \$5572.00. Previous research found that most fixed parts are "good as new," but 7 to 9 percent are not adequately repaired and will immediately fail again.³ Hence, we assume one gets $q=0.92$ for \$5572.00.

We also assume that 15 percent more expenditure per repair would solve the imperfect repair problem. Hence, if we model repair cost as aq^L , we find $a=6407.80$ and $L=1.68$. With this structure, $q=1$ repair costs \$6407.80 per unit repaired.

We also assume that a 50 percent modification (i.e., the modified unit will require removal with half the probability of a nonmodified unit) would entail \$5 million in fixed cost plus 10 percent of the relevant part's cost. Given our average part value, this per-unit $m=0.5$ cost is \$19,857.85. We assume that modification costs increase geometrically and that $m=1$ is prohibitively expensive.

Figure 7 shows the unit cost of a repair associated with the contractor's choices of q (solid line) and m (broken line).

We also assume that it is costless for a contractor to speed up the hands-on portion of repair. Hence, a contractor can choose any t_m or t_r without cost. We do, however, assume $t_m = t_r$, for simplicity. Of course, hands-on repair time is only a portion of total repair time as items may wait in repair queues if there is limited repair capacity.

²These data and the repair cost data come from the March 1993 issue of the Air Force Materiel Command's Recoverable Consumption Item Requirements System (D041) database.

³For example, Dumond, Eden, McIver, and Shulman (1994) argue that about 9 percent of F-15C/D radar units are "lemons."




Figure 7—Unit Cost of a Repair

AGGREGATION BIASES

In the real world, every aircraft has a large number of parts. In the application in this report, we aggregate aircraft parts such that

each aircraft has only one part. This aggregation procedure is not entirely benign. There are at least two resultant biases, and it turns out that they work in opposite directions.

Suppose, in the real world, one had an aircraft that experienced multiple, concurrent part failure. In the real world, two or more broken parts would enter the repair process.

In our aggregated world, each aircraft has only one part. Hence, this multiple, concurrent part failure on one aircraft would be represented by having one broken part enter the repair process. In this situation, our capacity estimates (x) are biased downward.

Another real world scenario, however, would have two different parts on two different aircraft fail. Two parts would enter the repair process. However, in the real world, cannibalization across aircraft may be possible. In this situation, only one aircraft may consequently be grounded. In the aggregated world, each aircraft with a broken part would be grounded. Hence, the contractor would likely buy extra repair capacity to avoid having to purchase extra spares, as compared to the real world situation with cannibalization.

Hence, there are countervailing biases introduced by the part aggregation procedure. If intra-aircraft part failures are highly correlated, the first bias (x too low) may be important. If, however, cannibalization is important in the real world, the second bias (x too high) may be important.

We ignore potential economies of scope in repairing multiple types of items. Also, we do not consider flexibility tradeoffs, e.g., flexible versus specialized repair technology.

Ultimately, however, the primary purpose of this report is to illustrate how contractors might respond to different types of contracts. It isn't clear there would be any bias in the comparative statics of contractor responses, even if the capacity estimates are biased in an unknown direction.

REFERENCES

- Abell, John B., and H. L. Shulman. *Evaluations of Alternative Maintenance Structures*. Santa Monica, Calif.: RAND, R-4205-AF. 1992.
- Commission on Roles and Missions of the Armed Forces. *Directions for Defense*. 1995.
- Dumond, John, Rick Eden, and John Folkeson. *Velocity Management: An Approach for Improving the Responsiveness and Efficiency of Army Logistics Processes*. Santa Monica, Calif.: RAND, DB-126-A. 1994.
- Dumond, John, Rick Eden, Douglas McIver, and Hyman Shulman. *Maturing Weapon Systems for Improved Availability at Lower Costs*. Santa Monica, Calif.: RAND, MR-338-A. 1994.
- Feinstein, Jonathan S., and Edward G. Keating. "An Economic Analysis of Life Care." National Bureau of Economic Research Working Paper No. 4155. September 1992.
- Feldman, Phillip. "Random Number Generators," in *Discrete-Event Simulation for Performance Evaluation of Systems*. New York: John Wiley & Sons (forthcoming).
- Hodges, James S. *Six (or So) Things You Can Do with a Bad Model*. Santa Monica, Calif.: RAND, N-3381-RC. 1991.
- Karlin, Samuel and Howard M. Taylor. *A First Course in Stochastic Processes*. New York: Academic Press, 1975.
- Leland, Jonathan W. *Core Methodology: A Critique and Generalization*. Alexandria, VA: Center for Naval Analyses, CRM 95-33. May 1995.
- Lewin, Arie Y., Kalman J. Cohen, and Richard C. Morey. *Decision Process Models of Contractor Behavior: The Development of Effective Contract Incentives (Phase II)*. Air Force Business Research Management Center. BRMC-81-5034. May 1984.
- Marsaglia, George, and Arif Zaman. "A New Class of Random Number Generators," in *The Annals of Applied Probability*, Vol. 1, No. 3, August 1991, pp. 462-480.
- McAfee, R. Preston, and John McMillan. *Incentives in Government Contracting*. Toronto: University of Toronto Press, 1988.
- Press, William H., Brian P. Flannery, Saul A. Teukolsky, and William T. Vetterling. *Numerical Recipes in C*. Cambridge: Cambridge University Press, 1988.

Press, William H., Brian P. Flannery, Saul A. Teukolsky, and William T. Vetterling. *Numerical Recipes in Pascal*. Cambridge: Cambridge University Press, 1989.

Schmalensee, Richard, and Robert D. Willig. *Handbook of Industrial Organization*. Amsterdam: North-Holland, 1989.

United States Air Force. *Defense Business Operations Fund. FY 1995 Budget Estimates. U.S. Air Force Overview. Operating/Capital Budget*. February 1994.

United States Air Force, Lean Logistics Office. *USAF Baseline Lean Logistics Master Plan and Road Map. Version 3.0. January 31, 1995*.