THE ECONOMICS OF ENTERPRISE TRANSFORMATION: AN ANALYSIS OF THE DEFENSE ACQUISITION SYSTEM

A Thesis
Presented to
The Academic Faculty

by

Michael J. Pennock

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
School of Industrial and Systems Engineering

Georgia Institute of Technology
May 2008

THE ECONOMICS OF ENTERPRISE TRANSFORMATION: AN ANALYSIS OF THE DEFENSE ACQUISITION SYSTEM

Approved by:

Dr. William Rouse, Advisor
School of Industrial and Systems
Engineering
Georgia Institute of Technology

Dr. Paul Griffin
School of Industrial and Systems
Engineering
Georgia Institute of Technology

Dr. Pinar Keskinocak
School of Industrial and Systems
Engineering
Georgia Institute of Technology

Dr. Stephen Cross
School of Industrial and Systems
Engineering
Georgia Institute of Technology

Dr. Kenneth Boff
Tennenbaum Institute
Georgia Institute of Technology

Date Approved: February 27, 2008

To my wife, Amanda.

ACKNOWLEDGEMENTS

In the course of producing this dissertation, there have been several individuals who helped facilitate its successful completion. First and foremost, I would to thank my advisor, Bill Rouse, for providing me the freedom to explore my ideas. The completion of this work was by no means a linear process. It involved several discarded thesis topics, a number of research dead ends, and the production of several working papers that did not end up in this dissertation. All throughout, Dr. Rouse provided enthusiastic support and advice that helped me endure the sometimes frustrating obstacles that arise in the pursuit of a Ph.D.

I would also like to thank all of the various individuals, some anonymous, who reviewed part or all of the material that comprises this dissertation and provided invaluable feedback. First, I would like to thank my committee members Ken Boff, Steve Cross, Paul Griffin, and Pinar Keskinocak. I would also like to thank Mike McGrath and Jack Gansler for providing feedback on the working paper that became Chapter 3.

I would like to thank my officemate and fellow Ph.D. student, Baabak Ashuri. Throughout these past four years he has served as a sounding board for my ideas, and on several occasions he provided critical input that saved me from some potentially embarrassing mistakes. Also, I would like to thank Diane Kollar who, as the office candy pusher, ensured that I had the energy required to complete this dissertation. I must also thank my dog who turned out to be the most reliable alarm clock I have ever owned. Thanks to him, I have not been able to sleep late once these past four years.

Finally and most importantly, I would like to thank my wife, Amanda. Not only was she willing to pick up and move so that I could pursue my Ph.D., but she also provided me with her love and support throughout this entire process.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

SUMMARY

Despite nearly 50 years of attempts at reform, the US defense acquisition system continues to deliver weapon systems over budget, behind schedule, and with performance shortfalls. A parade of commissions, panels, and oversight organizations have studied and restudied the problems of government acquisition with the objective of transforming the defense acquisition enterprise, yet the resulting legislative and procedural changes have yielded little, if any, benefit. Thus, the obvious question is why has acquisition reform failed? Three potential contributors were identified in the literature: misalignment of incentives, a lack of a systems view, and a lack of objective evaluation criteria. This dissertation attempts to address each of these problem areas.

First, I consider the issue of incentivization in the context of defense technology policy. A frequent criticism of defense acquisition programs is that they tend to employ risky, immature technology that increases the cost and duration of acquisition efforts. To combat this problem the Department of Defense rewrote their acquisition regulations to encourage a more evolutionary approach to system development. Nominally, this requires the use of mature technologies, but studies have revealed that acquisition programs continue to use immature technologies in spite of the new policies. To analyze this issue, the defense acquisition cycle was modeled as a stochastic process. Then, assuming that each acquisition program serves a diverse set of stakeholders, game theory was applied to show that the stable solution is to employ immature technology. It turns out that there is a tragedy of the commons at work in which the acquisition program serves as the common resource for each of the stakeholder groups to achieve its objectives. Since there is no cost to using

the resource, there is a tendency to overexploit it. The result is an outcome that is worse than if there had been a coordinated solution. Thus, the rational actions of stakeholders will lead to a contradiction of acquisition policy. Consequently, if the Department of Defense expects adherence to its evolutionary acquisition policy it must either strictly enforce technology maturity requirements or else realign incentives with desired outcomes.

Second, I evaluate cost and performance implications of the most recent defense acquisition transformation initiative, evolutionary acquisition. Proponents suggest that evolutionary acquisition will lower acquisition program costs, shorten delivery times, and improve the performance of fielded systems through the use of shorter and more incremental acquisition cycles. Supporting arguments focus on the impact of evolutionary acquisition on individual programs but fail to consider the defense acquisition enterprise as a system. To address this shortcoming, I analyze the impact of evolutionary policies through the use of a discrete event simulation of the entire defense acquisition system. It was found that while there should be an increase in the performance of fielded systems under evolutionary acquisition policies, the cost of operating the defense acquisition system as a whole does not inherently decrease. This is because the shorter acquisition cycles created by evolutionary polices mean that the overhead costs of each acquisition cycle are incurred more frequently. If these overhead costs do not decline sufficiently, the net cost to operate the acquisition system rises. This finding demonstrates the importance of considering the entire acquisition system before implementing a new policy.

Finally, I address the lack of objective evaluation criteria by developing a method to value acquisition process improvements monetarily. This is accomplished through the combination of price indices and options analysis. Since the US government is a non-profit entity, traditional cash flow based valuation methods are not applicable. Instead, the use of price indices captures the changes in the government's buying

power induced by acquisition reforms. This may be converted into an equivalent augmented budget stream that allows traditional investment evaluation tools to be applied. An additional advantage of the buying power method is that it captures the impact of the economies of scale inherent in the production of military systems. The augmented budget stream serves as the basis for applying options analysis, which properly accounts for the risk mitigating effects of staging. A comparison of this new method with more traditional methods reveals that only considering cost savings can significantly undervalue acquisition improvement opportunities, and even small improvements can have large returns.

CHAPTER I

INTRODUCTION

Despite nearly 50 years of attempts at reform, the US defense acquisition system continues to deliver weapon systems over budget, behind schedule, and with performance shortfalls. A parade of commissions, panels, and oversight organizations have studied and restudied the problems of government acquisition with the objective of transforming the defense acquisition enterprise. Despite some variation in the findings, several common threads have emerged. First, the Department of Defense (DoD) tends to pursue overly aggressive performance goals that require the use of immature technology. Maturation of technology can be quite unpredictable, and thus, early commitment to immature technology tends to lead to higher costs and longer development times. Second, acquisition decision making is decentralized, uninformed, and subject to the influence of multiple, diverse groups of stakeholders. This tends to lead to starting more programs than can be funded, duplicated work, and failure to consider potentially more cost effective alternatives. Third, program managers lack the authority and incentives to manage programs in a cost effective manner.

Pursuant to the identification of the aforementioned issues, acquisition reform panels and oversight groups typically make a number of recommendations to remediate acquisition failings. Often these recommendations are drawn from the study of private industry. The rationale is that market competition has honed the efficiency of private firms, and the DoD would do well to imitate their behavior. The Government Accountability Office (GAO), the investigative agency of Congress, has been particularly aggressive in pushing the reform of the defense acquisition system using commercial practices. More specifically, they recommend the strict enforcement of

technology maturity requirements; a staged, knowledge-based acquisition process; a centralized acquisition authority that manages the entire acquisition portfolio; and re-aligning the incentives, tenure, and authority of program managers with best interests of the acquisition system.

While the recommendations of the GAO and other reform entities are often nominally embraced by all relevant parties, including the DoD, there has really been little, if any, improvement. There are several key issues that contribute to this outcome. First, there is a failure of implementation. In many cases the DoD has reformed its acquisition policies only to allow programs to bypass them, or it implements only parts of a multipart recommendation resulting in ineffective outcomes. Second, the suggested reforms, while reasonable on the surface, are really just assertions. There is typically no systematic analysis of the impact of reforms in a defense context. Since there are important differences between the nature of the defense acquisition system and a commercial product development process, these fundamental structural differences may result in unexpected outcomes. Finally, commercial operations can reduce most decisions to a single objective, maximization of cash flow. This provides a universal standard of comparison for all policy alternatives. Government programs, however, are non-profit. They do not generate cash flows, and they attempt to satisfy multiple, non-commensurate objectives for multiple, diverse constituencies. Thus, there is no common measure on which to evaluate and compare acquisition policy alternatives. Consequently, a debate over the implementation of a policy reform essentially devolves into competing assertions without any objective means of resolving the situation.

This thesis attempts to address these three implementation issues through the use of systems engineering principles and economic theory, and one chapter is devoted to each of these issues. First, I consider the possibility that the underlying incentive structure of the defense acquisition system may lead participants to attempt to

bypass acquisition policies and recommended procedures. Second, I take a systems view of reform initiatives to evaluate their impacts within the context of the acquisition system. Third, I devise a valuation scheme that allows for the evaluation and comparison of acquisition process improvements.

To provide context for the first two issues, I have chosen a single acquisition reform effort to analyze, evolutionary acquisition. Evolutionary acquisition is an approach recently embraced by the DoD and committed to acquisition policy. Traditional acquisition programs attempt large leaps in capability in a single acquisition cycle. Hence, they are sometimes referred to as revolutionary acquisition programs. They do so by utilizing immature technology, and, as a result, often tend to take on the order of 10-20 years and incur significant costs. Evolutionary acquisition, on the other hand, sets more modest performance goals and utilizes more mature technology. Proponents argue that it will shorten cycle times, reduce cost, and increase the performance of fielded systems. Despite the fact that the DoD has embraced evolutionary acquisition as its preferred approach, a recent survey by the GAO has found that almost none of the DoD's current acquisition programs are evolutionary [32–34]. It turns out that programs are routinely exempted from the technology maturity requirements necessary to maintain an evolutionary acquisition system.

In Chapter 2, I consider why, if evolutionary acquisition is supposedly superior, DoD continues to pursue a revolutionary approach to acquisition? To address this issue, I model the defense acquisition cycle as a stochastic process. Then, assuming that each acquisition program serves a diverse set of stakeholders, I apply game theory and find that the stable solution is a revolutionary acquisition policy. It turns out that there is a tragedy of the commons at work in which the acquisition program serves as the common resource for each of the stakeholder groups to achieve its objectives. Since there is no cost to using the resource, there is a tendency to overexploit it. The result is an outcome that is worse than if there had been a coordinated solution. Thus,

evolutionary acquisition should theoretically provide superior system performance, but the rational actions of stakeholders will lead to a revolutionary acquisition policy. The policy implication is that if the DoD wishes to employ evolutionary acquisition, technology maturity requirements must be strictly enforced or else incentives must be realigned.

Chapter 3 evaluates the assertion that evolutionary acquisition reduces costs and increases performance. In theory, acquisition program costs should be lower under evolutionary acquisition because the use of mature technology reduces technology development costs. However, this assertion fails to consider the entire acquisition system. Unlike commercial firms, the DoD is effectively the developer, the manufacturer, and the sole customer of its product. Thus, the question is really whether evolutionary acquisition is still cost effective when total ownership of the entire system life-cycle is considered?

To answer this question, I developed a discrete event simulation that models both the acquisition system and the defense R&D system. I then consider as my key experimental variable the maturity at which a technology is transferred from the R&D system to an acquisition program. What I found is that the overall costs do not necessarily decrease under evolutionary acquisition, and this is primarily a result of the shorter acquisition cycles. Each acquisition program incurs large system development and production costs. When acquisition cycles become shorter, these costs are incurred more frequently. If these costs do not decline sufficiently under under evolutionary acquisition, the overall result is a higher acquisition system operating cost. Despite the potential for higher net operating costs, evolutionary acquisition still results in higher performance from fielded systems, and, thus, there is a direct trade-off between cost and performance. However, it is not necessary to vary the technology policy to achieve a particular cost/performance trade-off. If one inserts a delay interval between evolutionary acquisition cycles to artificially lengthen them,

one can achieve the full range of cost/performance combinations. Of course, the resulting gaps in production may lead to difficulties in the defense industrial base, but this is beyond the scope of this thesis. Thus, evolutionary acquisition may not be a free lunch, but it does create the opportunity to trade cost and performance rather than simply accepting an expensive and slow acquisition system.

Chapter 4 addresses the issue of valuing process improvements to the defense acquisition system. As mentioned previously, the government does not generate profits, and it serves a diverse constituency. Thus, it is difficult to employ traditional decision analysis tools to evaluate policy alternatives. A common approach is to employ cost savings as a figure of merit and utilize NPV analysis. However, this approach misrepresents value in two ways. First, it fails to consider that the market for defense systems is essentially a monopsony. As such, the per unit price that the government pays is heavily dependent upon economies of scale. Since most defense industries have excess productive capacity, they exhibit increasing returns to scale. Thus, when costs on a defense program rise, the government is forced to cut the size of the production run. This increases the per unit costs further, and the production quantity decreases even more. On the positive side, however, a decrease in program cost produces the opposite effect. As a result, nominal cost savings will understate the gain from a process improvement. Second, NPV analysis fails to consider the staged nature of most investments. Staging reduces downside risk exposure, and thus, an NPV analysis will understate the value of a risky investment.

To address the first issue, I develop a method to value a process improvement as a change in buying power through the use of a pricing index similar to those used to measure inflation. This allows for the monetary valuation of process improvements and facilitates addressing the second issue through options analysis. Options analysis appropriately considers contingencies in the implementation of a process improvement. Through the combination of these two approaches, I show that failure to

consider production economics and staging can significantly understate the value of a potential process improvement and could lead to an inappropriate rejection of the option. Finally, the method developed allows decision makers to objectively compare a portfolio of acquisition process improvements.

This dissertation is organized in the following manner. The remainder of this chapter discusses background and issues associated with defense acquisition reform. Chapter 2 is entitled "A Game Theoretic Analysis of Defense Acquisition Technology Policy" and reveals a tragedy of the commons at work in the defense acquisition system. Chapter 3, " A Systematic Analysis of the Cost and Performance Impact of Acquisition Technology Policy," presents the simulation study that reveals that evolutionary acquisition does increase performance, but it may also increase cost. Chapter 4, "A Method For Valuing Defense Acquisition Process Improvements," discusses the application of buying power and real options to value process improvements for the defense acquisition system. Finally, Chapter 5 summarizes the findings of this dissertation and discusses avenues for future research.

1.1 The Defense Acquisition Enterprise

The United States defense acquisition system is one of the largest and most complicated business enterprises in the world. The budget wielded by the Department of Defense is greater than the gross domestic product of many nations. Its nominal purpose is to develop and acquire systems for the US military, but like any public institution, it serves a diverse set of constituencies and purposes. Thus, the defense acquisition system differs from a commercial enterprise in several key aspects as noted by Cancian [9]:

- There is only one buyer.

- There are very few suppliers.

6

- The user is concerned with performance not price.

- Contracts are signed years before the product is available, and consequently, must be based on estimates for cost, performance, and schedule.

- Performance is difficult to judge and often subjective.

- The enterprise operates with public funds. The use of public funds is held to a different standard than private funds.

- Decision making power is diffused throughout the executive and legislative branches of government.

- Decisions are made under public scrutiny.

Furthermore, it attempts to satisfy a number of conflicting goals including maximize performance, minimize cost, minimize acquisition time, minimize risk, maximize program control, maintain jointness and interoperability, preserve the industrial base, maintain fairness and propriety, and advance national socioeconomic goals [9]. Consequently, the defense acquisition system faces challenges that no private enterprise ever would. This has made it impossible to operate at the same standards of efficiency that one would expect from a private firm.

Given these difficulties, legislation and acquisition regulations attempt to enforce a rational and transparent process that provides justification for the systems being acquired. Fox enumerates the basic steps in the process as follows [26]:

1. DoD identifies a security threat or defense operational mission.

2. DoD, usually with assistance from contractors, designs an engineering development program to meet the mission need and draws up an acquisition strategy and budget.

3. Congress authorizes and appropriates funds for the program.

4. The administration releases funds for the planned program.

5. DoD and interested contractors develop detailed technical approaches to the program.

6. DoD prepares a contract statement of work, with formal or informal assistance from contractors.

7. DoD issues requests for proposals to interested contractors and arranges pre-proposal conferences for bidders.

8. Contractors submit proposals to DoD, where they are evaluated.

9. DoD selects one contractor (or more), and the parties sign a contract for development of the weapon system.

10. The contractor begins work under the contract and each party initiates negotiated changes and modifications where required or deemed desirable.

11. The contractor delivers items to DoD for testing and evaluation while the work is in progress.

Following successful evaluation and approval by the relevant authorities, the system enters the production phase. To support this process, the DoD has established three overlapping systems: the Joint Capabilities Integration and Development System (JCIDS), the Planning, Programming, Budgeting, and Execution System (PPBE), and the Defense Acquisition System (DAS). JCIDS is used to identify military needs. PPBE is used to allocate resources, and DAS is for managing product development and procurement. To further complicate issues, each of these systems is operated by a different organization. JCIDS is managed by the Joint Requirements Oversight Council (JROC). PPBE is operated by the Office of the Secretary of Defense (OSD), and the DAS is run by the Undersecretary of Defense for Acquisitions, Technology, and Logistics (USD/AT&L).

The Defense Acquisition System is where the bulk of the work of developing and acquiring a system takes place. It is operated via the Defense Acquisition Management Framework (Figure 1.1). The framework divides the life-cycle of an acquisition effort into five phases: Concept Refinement, Technology Development, System Development, Production & Deployment, and Operations & Support. Ideally, warfighter needs are identified by the military services and proposed solution concepts are identified through the JCIDS process. Once the JROC approves a system concept, it moves into the Concept Refinement phase of the acquisition management framework. There are three decision points in the framework called milestones. At each milestone, a program must demonstrate it has met the requirements to move from one phase to the next. The purpose of these milestones is to provide decision makers with the opportunity to make an informed decision regarding the future of a program.
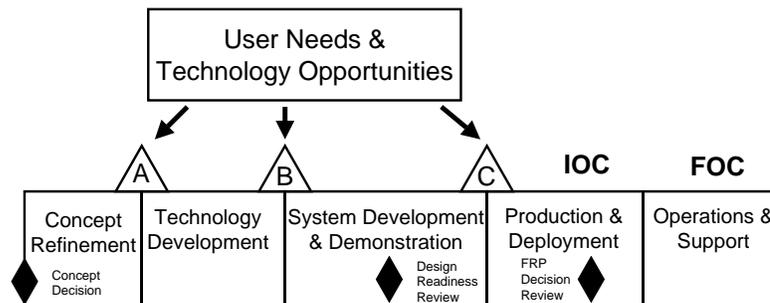
Figure 1.1: The Defense Acquisition Management Framework [19].

On the surface it would seem as if the DoD has established a rational and transparent means to acquire military systems. In practice, however, the system rarely operates as intended. Practically speaking, the military services (the Army, Air Force, Marines, and Navy) disproportionately influence the decision making process [41]. The services individually identify warfighter needs, and while these needs should be evaluated at the joint level, there are insufficient resources to adequately analyze service recommendations [36]. In the past, this has led to duplication and a lack of

interoperability. Furthermore, the JCIDS operates the continuously while the PPBE operates on a two-year cycle. This lack of synchronization means that OSD may not evaluate a proposed program for several years after the JROC's review. Thus, in practice, proposed programs are difficult to terminate following approval by the JROC because the sponsoring service begins budgeting and programming work immediately [36]. Consequently, it is difficult, if not impossible for OSD to manage or balance the acquisition portfolio through the PPBE process.

In yet another departure from the official process, acquisition programs often skip milestone requirements, and there is often significant concurrency between the phases of the defense acquisition framework [33]. Concurrency is used as a tool to shorten what would otherwise be a much longer acquisition cycle, but it is often at the price of increased cost and performance shortfalls. Beyond that, it has been suggested that concurrency works to advantage of the military services since it tends to shield programs from scrutiny until the system undergoes testing [41]. Once a program reaches this late stage, it is highly unlikely that it will be canceled.

The performance of the defense acquisition system has been decidedly mixed. Ultimately, the US military has acquired superior systems, but often well over budget and much later than expected. Cost overruns complicate budget allocation problems in several ways. Both defense contractors and the DoD systematically underestimate the cost to develop and acquire military systems [2, 10, 23]. This means that the DoD starts more programs than it can fund. Once a program begins to exceed its budget, authorities are forced to decide whether to underfund the program or reallocate funding from other programs. In particular, underfunding programs can take the form of stretching the program out over a longer time period and can lead to a higher total cost, or performance requirements may be loosened or dropped. This leads to funding instability that complicates efficient program management as well as delay or loss of anticipated military capability. The result is that it is difficult,

if not impossible, to rationally allocate the budget to meet military objectives, and the acquisition system is often criticized for producing a portfolio of weapons systems that does not meet national military objectives [41].

Schedule slippage also leads to poor options for decision makers. If a delinquent program is allowed to continue, warfighters will have to make due with their current and possibly obsolete equipment for longer than anticipated. Alternatively, performance objectives could be sacrificed in order to field the system faster. But that, once again, leaves warfighters with less capability than anticipated.

Because of these persistent problems with defense acquisition, transforming the defense acquisition enterprise has been a perennial objective of both the Presidency and the Congress for nearly fifty years. There have been a number of commissions, panels, and studies that have attempted to ascertain the cause of these problems and make recommendations to remedy the situation. The next section describes the history and substance of the efforts to transform defense acquisition.

1.2 Transforming the Acquisition Enterprise

Defense acquisition reform as it is thought of today began with the start of the Cold War. Persistent problems with inaccurate cost estimates and schedule slippage led Defense Secretary Robert McNamara to spearhead a long list of reforms during the Kennedy Administration. McNamara's initiatives fell into three broad groupings: program planning, source selection and contracting, and program management [66]. The purpose of these initiatives were to bring systems analysis and industrial practices to defense acquisition as well as multi-year planning and budgeting and an objective means of selecting contractors. The innovations introduced by McNamara yielded mixed results, and acquisition programs continued to exhibit the usual cost, schedule and performance problems.

Since McNamara there have been a host of initiatives and regulations designed to

Table 1.1: Timeline of Acquisition Reform Efforts. (Compiled from Christensen, et al., 1999 [10] and Rogers and Birmingham, 2004 [61].)

Year	Initiative, Legislation, or Regulation
1969	Packard Initiatives
1971	Blue Ribbon Defense Panel (Fitzhugh Commission)
1972	DoD Directive 5000.1, Commission on Government Procurement
1973	DoD Directives 5000.3 and 5000.4
1975	DoD Instruction 5000.2, DoD Directive 5000.28
1976	OMB Circular A-109
1978	Defense Science Board Acquisition Cycle Task Force
1979	Defense Resource Management Study
1981	Carlucci Initiatives, Defense Acquisition Improvement Program
1982	Nunn-McCurdy Amendment
1983	Grace Commission
1985	DoD Directive 5000.43
1986	Packard Commission
1987	DoD Directive 5134.1 and 5000.49
1989	Defense Management Review
1991	Revision of DoD Instruction 5000.2
1993	National Performance Review
1994	Federal Acquisition Streamlining Act, Perry Memo
1995	Federal Acquisition Improvement Act
1997	DRI Report
2000	The Road Ahead
2001	Rumsfeld's Transformation Vision
2002	Cancellation of 5000 series regulations
2003	Revision of DoD Directive 5000.1 and DoD Instruction 5000.2

improve defense acquisition. Table 1.1 provides a timeline of these acquisition reform efforts. One of the most notable attempts at acquisition reform was the Packard Commission report issued in 1986 [59]. The recommendations made by the Packard commission centered around such concepts as streamlining, better planning, better retention of quality personnel, and the adoption of commercial best practices. Despite attempts to implement the recommendations of the Packard Commission, a study on cost overruns found that changes made based on the Packard Commission report did not result in any reductions [10]. In fact, the same study revealed that from 1960 to 1990, there has been essentially no improvement in cost overruns.

A constant voice in the push for acquisition reform has been the Government Accountability Office (GAO). For years they have analyzed the performance of the defense acquisition enterprise and made recommendations for improvement. Most recently, they have pushed for the adoption of an acquisition enterprise comparable to a commercial new product development process. Commercial product development cycles are much shorter and more incremental that those found in defense acquisition. New commercial products tend to be evolutionary while new military systems tend to be revolutionary. The thought is that if the DoD adopted shorter, more incremental acquisition cycles, they would reduce risk, which would in turn reduce cost overruns and schedule slippage. Furthermore, warfighters would receive up-to-date equipment more frequently. This would make the entire system more adaptable to the ever changing threats to US national security.

To implement evolutionary acquisition, the GAO has suggested a number of commercial best practices including centralized management of the acquisition portfolio, a staged and centrally managed technology development process, technology maturity requirements, strict enforcement of milestones with knowledge-based requirements, reduction of concurrency between phases, and an improved incentive and retention structure for program managers. The DoD has taken steps to implement evolutionary acquisition. In particular, acquisition regulations DoD Directive 5000.1 [18] and DoD Instruction 5000.2 [19] were revised in 2003 to make evolutionary acquisition the preferred approach. However, the GAO has found that even under the new regulations, most defense programs are still revolutionary rather than evolutionary, and they continue to experience cost overruns and schedule slippage [33].

Thus, thus the ultimate question is why has nearly 50 years of attempts at reform yielded little or no improvement? Regarding the most recent initiatives, the GAO feels that the DoD has not sufficiently implemented commercial best practices to effectively transform the way it acquires systems [30–32, 34, 36]. Implementation has

certainly always been an issue. During the Reagan adminstration Defense Secretary Casper Weinberger and Deputy Secretary Frank Carlucci pushed a comprehensive set of acquisition reforms, but the GAO found that despite the strong support from upper level leadership, these, too, were never fully implemented [29].

Why does implementation falter? Table 1.1 would seem to indicate that it is not from a lack of initiative. In his paper "Acquisition Reform: It's Not as Easy as It Seems" [9], Cancian points out that the defense acquisition system represents a series of trade-offs among diverse constituencies, each with differing objectives. One person's waste is another person's essential requirement. Even minor reforms will create winners and losers, and losers may actively attempt to thwart reforms. Thus, there are often disincentives to follow through on acquisition reforms. Chapter 2 of this dissertation analyzes this aspect of implementation failure.

Furthermore, Cancian discusses the cyclical nature of acquisition reforms. They tend to vacillate between extremes. For example, a commission will find that excessive regulation and oversight leads to an additional cost burden for the government. Regulations will be altered to "streamline" the defense acquisition system. After a few years of operating under these rules several major programs will post major cost overruns because they were not adequately monitored and controlled. Of course, there is an outcry and a demand for more oversight, and the cycle repeats. Cancian has noted this behavior with such acquisition issues as cost-plus versus fixed-price contracts, system testing requirements, and the use of MILSPEC versus COTS equipment. Ultimately, this vacillation is related to a failure to understand and acknowledge the tradeoffs inherent to any changes in a system as complex as the defense acquisition enterprise. For example, using Commercial Off The Shelf (COTS) equipment instead of equipment built to military standards may reduce acquisition costs, but in some cases it will lead to reliability issues since many commercial products are not built to the demanding standards that military field use requires. This is not to suggest

that there is no place for COTS equipment in the DoD. Instead, the question is what is the appropriate balance. This viewpoint is further supported by Hanks, et al. [38] whose interviews with program mangers reveal that the current trend in streamlining and employing performance based approaches has gone too far and leads to vague contracts that are difficult to manage. Unfortunately, policy makers do not fully understand these trade-offs when they alter acquisition regulations, and thus, there is a tendency to cycle between the extremes. Chapter 3 of this dissertation attempts to provide a better understanding of these tradeoffs for the most recent acquisition initiative, evolutionary acquisition.

Finally, Cancian also notes that there are "few objective criteria by which to judge defense activities and outcomes" [9]. Consequently, it is difficult to apply commercial practices for evaluation and control to defense acquisition. Furthermore, it complicates the comparison of acquisition policy alternatives since much of the perceived benefit of military systems is subjective. Chapter 4 of this dissertation provides a partial solution to this problem by providing a method for valuing acquisition process improvements monetarily.

1.3   Acquisition Research

Given the size, complexity, and level of import of the defense acquisition system, there has been surprising little academic research on the topic. Some work is being performed by government affiliated organizations such as internal DoD organizations and the Federally Funded Research and Development Corporations (FFRDCs), but it tends to be limited in scope [27]. Examples include topics such as whether or not to use lean manufacturing during military aircraft acquisition [13] or whether competitive sourcing is a cost effective means of purchasing services [60]. While these topics are important, there is little work on the acquisition enterprise itself and the policies and organizations that drive it. As Gansler and Lucyshyn state, the work

being performed is "not a substitute for dedicated, replicable academic research" [27].

To address this research vacuum, the Naval Postgraduate School has begun an Acquisition Research Program to sponsor research at academic universities, and in 2004 it began hosting an annual acquisition research symposium. However, these initiatives are still in their infancy. A survey of the conference proceedings revealed that the vast majority of the work is still limited in scope and supported by anecdotal evidence [57]. Even so, the quality of work is improving. A system dynamics analysis of spiral development presented by Dillard and Ford [16] is one such example of how academic research methods can be applied to defense acquisition to achieve informed policy recommendations. Even so, little research has been done on the acquisition enterprise as a system.

Two notable exceptions are a paper by Clark, Whittenberg, and Woodruff [11] and a PhD dissertation by Burgess [8]. Both works take a systems dynamics approach to modeling the complex macroscopic interactions between the defense acquisition system, the US economy, the US government, and the threat posed by the Soviet Union. While both are excellent examples of taking a systems approach to studying defense acquisition, neither attempts to perform policy analysis with their respective models. Thus, they make no real recommendations as to how to improve defense acquisition. Furthermore, the models that they developed are so high-level that it may be difficult to link the impact of actionable policies to model components without expanding the models to include the behavior of lower level sub-systems.

While most research has not taken a systems approach, there have been some noteworthy academic works that cover certain aspects of the acquisition problem. The economics of the defense industry, in particular, has drawn some attention from the academic community. Notable works include Peck and Scherer's economic analysis of the structure of the defense industry and acquisition process [56], Scherer's analysis of contract incentives between the government and defense contractors [68],

and Baldwin's analysis of the market features of the defense industry [3]. In particular, Peck and Scherer's work correlated with the early attempts at acquisition reform championed by Robert McNamara that were discussed above. Fox considered his widely cited work *Arming America* to be a successor to Peck and Scherer's work, and he focused on the management of the acquisition process [25]. Weida and Gertcher considered the interaction of politics and economics in national defense and in particular why defense decisions may not be economically efficient [74]. Gansler provides an extensive analysis of the state of the defense industry in the post-Vietnam era [28]. In particular, he emphasized the importance of intelligently managing the defense industry because of the lack of the traditional market forces that ensure economic efficiency. Finally, Sandler and Hartley provide a broad coverage of the application of macro and microeconomic theory to the concept of national defense covering issues such as alliances and arms races, the industrial base, disarmament, and the arms trade [67].

Much of this previous work deals with the unusual industrial economics of producing military systems, and, in particular, the interaction between government and industry. One common issue is the way in which defense contracts are written and administered impacts the efficiency of the defense industry. Scherer extensively considers the nature of contract incentives, and Sandler and Hartley note that game theory has been applied to interactions between contractors and the military.

Thus, while there has been some academic research on the topic of defense acquisition, there has not been nearly enough considering the size and import of the problem. The primary output of past work has been recommendations to improve or reform the defense acquisition system. Perhaps the most important difference between this dissertation and past academic work is that it considers why reforms have failed at the implementation stage. In particular, this dissertation takes the approach that the defense enterprise is a system that can be analyzed to better understand the

implementation issues surrounding defense acquisition reform. Consideration of how the components of the enterprise interact is necessary to understand why a policy, however well-intentioned, fails to realize its intended objective. By treating an enterprise as a system [62], the entire suite of approaches and tools developed to support systems engineering, industrial engineering, and operations research is available to analyze the impact of changes to the defense acquisition system. Such a systematic approach also provides the opportunity to adapt economic principles and investment analysis to a domain where, traditionally, they have been considered inapplicable. Thus, while many possible causes of implementation failure have been postulated, the methods developed in this dissertation provide a means to analyze them.

CHAPTER II

A GAME THEORETIC ANALYSIS OF DEFENSE ACQUISITION TECHNOLOGY POLICY

Chapter 1 asserted that participants in the defense acquisition system may be disincentivized to implement or follow proposed acquisition reforms. This chapter tests that assertion regarding one of the most contentious issues of defense acquisition, technology maturity policy. As was mentioned previously, the consistent use of immature technology in acquisition programs has often been cited as a key driver of cost overruns, schedule slippage, and performance shortfalls. But if the use of immature technology is so widely recognized as a causal factor in the failure of defense acquisition, why does the practice continue?

There have been repeated calls for the Department of Defense to use evolutionary rather than revolutionary acquisition strategies. As mentioned previously, evolutionary acquisition relies on shorter acquisition cycles and mature technologies to achieve more modest goals. Supporters believe that this approach will lead to an overall improvement in the level of capability deployed as well as lower cost acquisition programs.

In fact, the DoD has revised its acquisition polices to that end [30]. Despite these new policies, recent Government Accountability Office (GAO) reports have indicated that most major acquisition programs are still revolutionary rather than evolutionary and do not follow current DoD guidelines for knowledge-based acquisition [32–34]. It seems that every program is an exception. Why is this?

Two questions logically follow: What level of maturity is acceptable for a technology to be included in a major acquisition program, and what has been preventing

the DoD from practicing an evolutionary acquisition process?

The analysis presented in this chapter will show that the answers to these questions are more broad than the evolutionary vs. revolutionary acquisition debate. It turns out that given the risks and structure of each acquisition program, there is a performance optimal technology policy that will maximize the gain in capability over time. However, when acquisition programs attempt to serve a diverse set of stakeholders, a tragedy of the commons arises where each stakeholder is incentivized to push for technology that is more immature than is optimal. The result is that all stakeholders in the acquisition program end up worse off.

Thus, there is a natural tendency towards revolutionary acquisition, and any acquisition strategy that advocates the use of more mature technology is inherently unstable. The policy implication for the US Department of Defense is clear; either technology maturity requirements must be strictly enforced or incentives must be realigned.

2.1 Background

A common criticism of the defense acquisition process is that it tends to emphasize large leaps in capability achieved by utilizing promising but immature technology. Changes to defense acquisition policy over the last several years have attempted to reverse this trend by creating a milestone process where programs must meet certain requirements before proceeding from one phase to the next [18, 19] (See Figure 1.1). Part of this milestone process is an assessment of the maturity of technologies to be employed in an acquisition program as well as a plan to manage their development.

Technological maturity is typically assessed using the Technology Readiness Level (TRL) scale (See Appendix E). The TRL scale is a qualitative assessment scale that is designed to aid decision makers by providing some sense of a given technology's

level of risk. In general, one assumes that the higher the TRL level, the less uncertainty a technology brings to a program. It is important to note that the TRL scale evaluates a technology in isolation and does not consider the integration risks [71]. Regardless, the aforementioned policy changes encourage programs to utilize more mature, demonstrated technologies (i.e., higher TRL levels) rather than more immature and consequently, more risky technologies. For example, entering the system development phase nominally requires all critical technologies to be at TRL level 6 or higher (Though the GAO recommends at least TRL level 7 [19]).

What is the rationale behind a policy that requires a relatively mature level of technology? The issue is that development of immature technology is fairly unpredictable in terms of cost, schedule, and efficacy. When a program contains multiple immature technologies, it tends to delay the program and add cost. If technology development is done in concurrence with system development, the problem can be exacerbated because unforeseen outcomes can lead to significant rework. The net result is that, on average, programs with immature technologies will take longer and cost more. Consequently, warfighters must make due with obsolete equipment longer and, thus, increasing the chances that they will engage in combat operations with less capability than they could have had otherwise.

As a result, it would seem that a superior approach would be to reduce cycle time by setting more modest goals for each deployed increment of capability. This is what the GAO refers to as an evolutionary rather than a revolutionary acquisition process, and there are several ways to achieve such a process. First, one can make use of open-architecture design and spiral development. The idea behind spiral development is that the system can be deployed with an initial mature technology, which can then be upgraded over time [43]. This approach can work well for technologies that are loosely coupled to the system design. In other words, there is a clear, well-defined interface such that changes in the implementation of the subsystem or technology to

be upgraded do not interfere with the rest of the system. Open architecture design is perfect for a technology such as a software algorithm. Assuming that the software interface has been standardized, it is comparatively straightforward to replace an old software component with a new one. This approach, in fact, has been demonstrated successfully on submarine acoustic systems [7].

When technologies or subsystems are tightly coupled to the overall system, however, any changes to the design of the subsystem will impact the design of the whole system [12, 24]. Thus, spiral development is not always a feasible alternative. An extreme example would be the hull-form of a surface combatant. If some critical issues were to arise with the hull design, it is likely that a significant portion of the ship would have to be redesigned. Of course, hull form is a rather obvious case, but there are many mission critical systems in any modern military system that exhibit varying degrees of interaction with the rest of the system design. Since changes to these systems would require substantial rework, it is imperative that they be mature prior to system integration, hence the appeal of evolutionary acquisition.

Under evolutionary acquisition, system acquisition cycles are more rapid and make use of mature, available technology. The development of new technologies is detached from the acquisition process, so that the fate of a program does not hinge on the success or failure of any one risky technology. The evolutionary approach is enforced via a knowledge based acquisition process. The program contains a number of evaluation points or milestones. At each milestone the program must demonstrate that it has met certain developmental requirements in order to proceed to the next phase. For example, Milestone A entails requirements such as an Initial Capabilities Document, an Analysis of Alternatives (AoA), a Systems Engineering Plan (SEP), and Technology Readiness Assessment.

Despite the fact that the DoD acknowledges evolutionary and knowledge-based acquisition as best practices and has committed them to policy, recent GAO reports

have indicated that most major acquisition programs do not follow these polices [32–34]. Consequently, these major acquisition programs have continued to experience significant cost overruns and major delays. In particular, these reports have indicated that most major acquisition programs are revolutionary rather than evolutionary, and they are permitted to bypass major milestone requirements. Most rely on multiple immature technologies that are not fully developed before overall system development begins. The Office of the Secretary of Defense (OSD) has acknowledged that this is a common practice [33].

One example in particular that makes the consequences of this acquisition approach clear is the case of WIN-T and JNN-N. The Warfighter Information Network-Tactical (WIN-T) is the next generation tactical communications network for the US Army and will provide a major leap forward in battlefield communications. However, when the program moved into the system development phase, 9 of the system's 12 critical technologies were immature [34]. As a result, WIN-T has been unavailable for both Operation Iraqi Freedom and Operation Enduring Freedom. Because it was determined that there was an urgent need for better battlefield communications to support these two operations, the Joint Network Node-Network (JNN-N) program was created. To address this urgent need, the JNN-N program bypassed many of the normal acquisition procedures to accelerate fielding of the system. While this may be understandable given the urgency of the situation, acquisition procedures are in place to ensure that acquired systems function properly and are cost effective. As the GAO points out,

> When the Army opted to pursue large technology advances in networking capabilities to support the future forces through WIN-T, rather than pursuing a more incremental approach, it accepted a gap in providing tactical networking capabilities to the warfighter ... If the Army had followed

DOD's acquisition policy preferences, which emphasize achieving capabilities in increments based on mature technologies to get capabilities into the hands of the user more quickly, it might have been able to get needed communications capabilities to the warfighter sooner. [34]

Thus, a more evolutionary approach to acquisition may have reduced the risks to the warfighter by avoiding capability gaps as well as mitigating the need for emergency programs that bypass the usual acquisition procedures.

To summarize, the Department of Defense claims to favor evolutionary acquisition, but does not follow through in practice. The GAO asserts that there are a number of causes, one of which is the lack of mandatory controls on the milestone process [30,32–34]. But if evolutionary acquisition is superior, why would the DoD not follow its tenets even without the mandatory controls? Or to put it more broadly, why does the DoD appear to be working against its own best interests?

2.2 Modeling Approach

To address the questions at hand, we must reduce an acquisition program to its most basic and essential features. The objective here is not to produce a predictive model, but instead to better understand the underlying forces that drive the behavior of participants in the acquisition system. To that end, this analysis is predicated on three assumptions about the defense acquisition system.

1. Acquisition programs are dependent upon the development of multiple critical technologies.

2. There are multiple parties that have differing interests with regard to the outcome of an acquisition program.

3. The requirements process for acquisition programs is negotiable, or at least open to influence from the aforementioned parties.

Anyone familiar with the defense acquisition system will likely find these assumptions apt. The underlying hypothesis of this chapter is that because acquisition programs are government programs, they are subject to the influence of multiple groups of stakeholders. These stakeholders may have different demands for the capabilities provided by an acquired system. This influences the technology policy of the acquisition program and, consequently, the cost and duration of the program. When the requirements process is not tightly regulated, competing demands from stakeholders will lead to an overly aggressive technology policy that results in longer acquisition programs and lower fielded system performance over the long-run. Essentially, this a tragedy of the commons where the acquisition program serves as the common resource.

To investigate this hypothesis, we must develop a model of an acquisition program, and if we wish to find the basic underlying forces that drive acquisition technology policy, we must reduce an acquisition program to its most essential features. In that vein, several simplifying assumptions are required.

First, it is assumed that each acquisition program consists of two phases: a technology development phase and an integration phase. In the technology development phase, critical technologies are matured to the point that they can be utilized in the final system. Once this phase is complete, system integration can move forward. It is further assumed that each critical technology can be developed in parallel, but all must be complete before system integration can begin. This is an admitted simplification that works both for and against the acquisition program. The assumption of parallel technology development is somewhat optimistic as the outcome of each critical technology may be somewhat interdependent. The assumption that all development must be completed is somewhat pessimistic because some integration work can be done based on the estimated outcome of technology development. However, in should be noted that since unanticipated outcomes in the technology development

phase can lead to substantial rework in the integration phase, too much concurrency can undermine any time savings. Ultimately, there is a limit to how much time can be saved through phase concurrency. Thus, if we were to weaken the no concurrency assumption, it would serve to somewhat dampen the impact of development risk, but it would not materially change the results of this analysis.

Given these assumptions we can structure each acquisition program as shown in Figure 2.1. The purpose of each acquisition program is to improve upon a set of
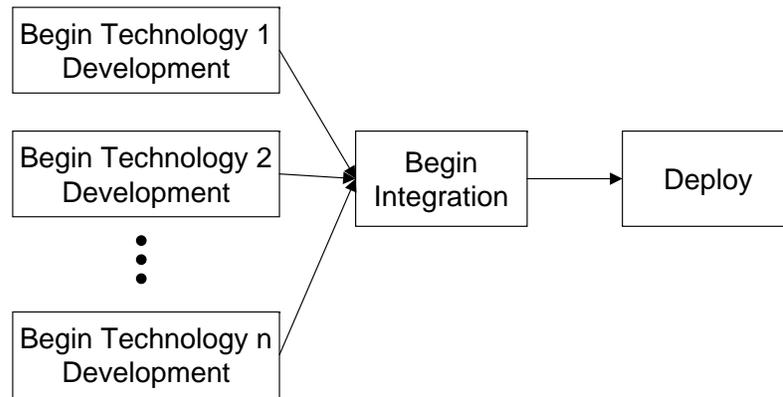
Figure 2.1: Conceptual model of a defense acquisition program.

capabilities provided by one or more existing systems. While this model treats capabilities in an abstract sense, notionally one could consider performance in areas such as target detection, survivability, and lethality as examples of capabilities provided by a deployed system. So the presumption is that a successfully completed acquisition program would improve upon what is currently provided. To account for performance improvements, the metric of interest will be the percent improvement of each capability an acquired system provides over what is currently fielded.

It is assumed that each technology development activity is independent, and there is a one-to-one mapping between technology development activities and the set of capabilities provided by an acquired system. (It is assumed that if multiple development activities are required to achieve a particular capability, these are consolidated

into a single activity.) With this linkage established, it is clear that the technologies selected will have an impact on the level of capability improvement achieved. Presumably, larger improvements in capability require more immature technologies. Unfortunately, immature technologies typically take longer to develop and entail more risk. Thus, the technology development activities are inherently stochastic.

We can model the time to complete an acquisition program as

$$P = \max(X_1, X_2, \ldots, X_n) + I \tag{2.1}$$

where P is the time to complete the acquisition program, X_i is the time to complete technology development activity i, and I is the time required for the integration phase. Since technology development is stochastic, each X_i is a random variable governed by a non-negative distribution function. From Equation (2.1) it is immediately apparent that the maximization of several random variables will drive the behavior of this model. Any stochastic behavior of I will not materially affect the results of this analysis. So, for simplicity, it is assumed that I is deterministic.

The next crucial feature of the model is the linkage between the distribution of each X_i and the maturity of the technology selected. Keeping in line with the previous discussion, it is assumed that there is a relationship between the percent increase in capability provided by the system to be acquired and the time required to develop the requisite technologies. In particular, for this analysis, the desired level of capability improvement is always achieved, but the time required is uncertain. More specifically, it is assumed that the expected time to complete a development activity increases as the required capability improvement increases. Furthermore, it is assumed that it requires zero time to achieve zero improvement.

But there is really more to this relationship. It is certainly not linear. In fact, one would expect a diminishing return to scale. In other words, there is some benefit to developing technologies in steps since one learns from each step. Similarly, we would also expect there to be less risk with a smaller step. Thus, there is a price to pay for

attempting one large leap in capability all at once. Graphically, we would expect a relationship such as that depicted in Figure 2.2.

Expected Time for Technology Development
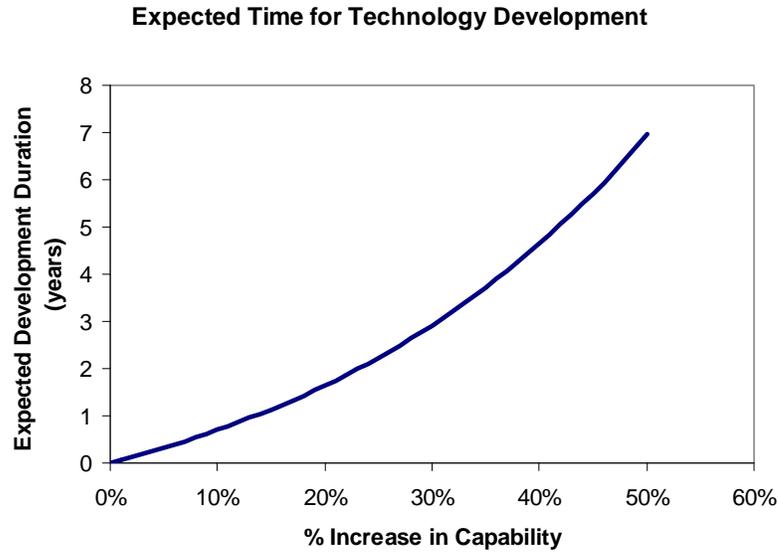
Figure 2.2: Notional relationship between the targeted percent increase in capability versus the expected time to develop the required technology.

Of course, when a program involves the development of multiple technologies to improve multiple capabilities, the expected duration of the program will depend on the distribution of the maximum of all development efforts. Because of this interaction effect, one must consider the level of improvement desired for each provided capability. Thus, a technology policy for an acquisition program consists of the targeted level of improvement for each of the capabilities to be provided by the acquired system.

Unfortunately, considering one acquisition program in isolation is not particularly useful. Instead, it is the long-term performance of a technology policy that is of interest. There is some notion that there is an optimal target for capability improvement. Too low, and we loose too much time due to the overhead inherent in an acquisition program. Too high, and time is wasted chasing overly difficult or technically infeasible approaches to improving capability. This behavior is only realized, however, when we consider the compounding effects of a sequence of programs. To that end, it is

assumed that as one acquisition program completes, another is begun immediately to provide the next upgrade in capability. This yields a stair-step capability trajectory as depicted in Figure 2.3.

Deployed Capability Trajectory
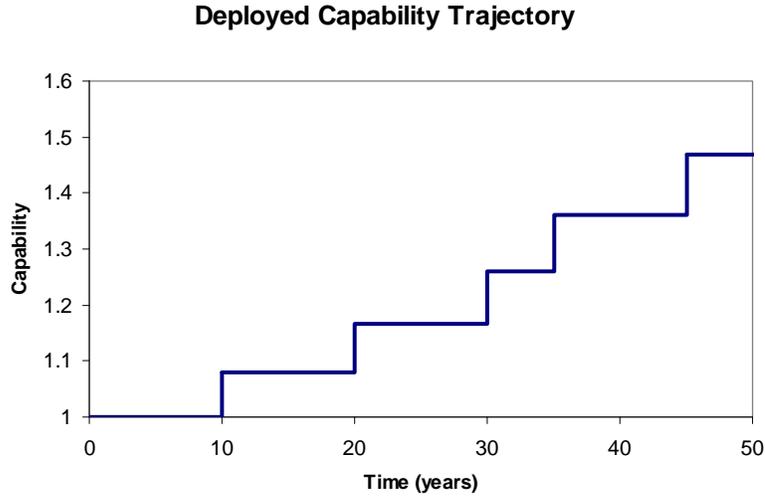
Figure 2.3: Sample capability trajectory.

To measure the performance of a particular technology policy, say a capability improvement of 10% for each acquisition cycle, we consider the long-run effective annual capability growth rate. To better understand this metric, consider the accumulation of interest in a bank account. For example, assume that a bank account pays 1.25% interest every quarter. With the compounding effect, that is equivalent to receiving 5.1% annually. The effective capability growth rate is analogous. It translates the sequence of discrete improvements in capability provided by acquisition programs into an equivalent annual rate of growth. Since the outcome of the programs is stochastic, however, we are concerned with the long-run behavior of this metric.

There is one final point to note regarding this modeling approach. That is that the compounding effect of repeated technology development efforts implies a perpetual exponential increase in the level of capability provided by acquired systems (Something akin to Moore's law for microchips). This is admittedly an assumption but, a

reasonable one. While particular technologies certainly reach points of diminishing returns, a particular capability may be provided by a variety of technologies. For example, when aircraft speeds became limited by piston engine technology, switching to jet propulsion allowed speeds to continue to rise. Of course, no capability can continue to improve indefinitely, but for the time horizons that most decision makers in government would consider, the assumption of compounded growth is not unreasonable.

That being said, in the following section we will translate the conceptual model described above into a mathematical model. That will allow us to better understand what factors influence technology policies in defense acquisition.

2.3 Analysis

In order to construct our model, we must first define some notation. The most basic description of the model was expressed in Equation 2.1, except now we recognize that program duration is dependent upon the number of technologies and the technology policy. To that end, we define the following:

X_i is the duration of development for capability i.

$g_i \geq 0$ is the targeted percent increase in capability i.

n is the number of capabilities provided by the system to be acquired.

$X_M = \max(X_1, X_2, \ldots, X_n)$ is the duration of the technology development phase where each X_i is independent.

$P = X_M + I$ where I is an exogenous integration time, and P is the total program time.

$X_i \sim F_i(x; g_i)$ where $F_i(x; g_i)$ is a non-negative, continuous, and differentiable probability distribution function and g_i is a parameter.

It is a well known result that the distribution of the maximum of several indepen-
dent random variables is the product of the respective distribution functions.

$$X_M \sim F_1(x)F_2(x)\cdots F_n(x)$$

Thus, we can readily describe the distribution of program duration, P. To capture the
behavior of technology development activities with regard to the selected technology
policy, we define the following:

Let $\boldsymbol{g}$ be the vector of the n capability targets, g_i.

Let $Y(\boldsymbol{g}, n) = E[P] = E[X_M] + I$ be the expected duration of the program.

Let $W_i(g_i) = E[X_i]$ be the expected time to develop capability i.

We must also make some assumptions about the behavior of $W_i(g_i)$, and how
it relates to the distribution of X_i. First, we assume $W_i(0) = 0$, $W_i'(g_i) > 0$, and
$W_i''(g_i) \geq 0$ for $g_i \geq 0$. These assumptions assure that it takes zero time to do
nothing, that the expected activity duration strictly increases as a larger leap in
capability is attempted, and that the expected duration increases at increasing rate.
All of these assumptions are consistent with the discussion in Section 2.2.

Second, we must define how changing g affects $F_i(x; g)$. This is accomplished by
utilizing an affine transformation on the random variable. If we have a particular
non-negative random variable, X, its expectation can be altered through the affine
transformation $Y = aX + b$. A nonzero value for b essentially translates the distri-
bution of X but does not change the spread. As mentioned previously, however, we
would expect the risk to increase with g, so a pure translation would not be appropri-
ate for this analysis. If a is not equal to one, then the distribution will be stretched
asymmetrically. Thus, both the expected value and the spread of the distribution will
change. For this analysis, it is assumed that $b = 0$ to both simplify the analysis and

leave open the possibility that a technological breakthrough, though unlikely, would allow development to proceed quickly even for an aggressive capability goal.

Returning to our model, if we know the distribution of X_i for some particular value of g, say g', we can determine the distribution of X_i for every other value of g. Let us define the random variable D such that $D = X_i$ when $g = g'$. Thus, $X_i = a(g)D$ where the affine transformation is a function of g. If we define g' such that $W_i(g') = 1$, then we know that $E[D] = 1$, and consequently, $a(g) = W_i(g)$. Thus, our affine transformation becomes $X_i = W_i(g)D$ and has the requisite property $E[X_i] = W_i(g_i)$. In this manner we can relate the change in g to the change in the distribution of a technology development activity. While this approach to defining the relationship between the distribution of X_i and g may seem arbitrary, it is equivalent to changing the λ parameter for an exponential distribution or the μ parameter for a lognormal distribution. Thus, it is merely a generalization of the impact of changing the parameters for several common distributions and avoids tying the results of the analysis to any one distribution type.

As stated previously, the objective of this analysis is to find the long-run effective annual capability growth rate for a given technology policy. The technology policy is determined by selecting a targeted percent increase in each capability over the next acquisition cycle. Since each acquired platform may provide more than one capability, outcomes of the realized capabilities are interdependent.

In this model, acquisition programs occur in sequence with uncertain durations. This constitutes a renewal process where P is the inter-arrival time and $N(t)$ is the number or arrivals (i.e., completed programs) at time t. In this framework, the effective annual growth for a given capability is

$$(1 + g_i)^{\frac{N(t)}{t}} - 1.$$

Note that $\frac{N(t)}{t}$ is the annual arrival rate (assuming t is in years). Of course, we are interested in the long-run, so by the strong law of the renewal process,

$$\frac{N(t)}{t} \to \frac{1}{E[P]} \text{ as } t \to \infty.$$

Therefore, the long-run effective annual capability growth rate is

$$(1 + g_i)^{\frac{1}{E[P]}} - 1.$$

Let $V(\boldsymbol{g}, n) = \frac{1}{E[P]}$ where $\boldsymbol{g}$ is the vector of technology policies. Then the optimal policy is

$$\max_{\boldsymbol{g}} \quad (1 + g_1)^{V(\boldsymbol{g},n)} - 1, (1 + g_2)^{V(\boldsymbol{g},n)} - 1, \ldots, (1 + g_n)^{V(\boldsymbol{g},n)} - 1.$$

Thus, we are faced with a multi-objective optimization problem. If we assume that all technologies are symmetric (i.e., $W_1(g_1) = W_2(g_2) = \cdots = W_n(g_n) = W(g)$ and $F_1(x) = F_2(x) = \cdots = F_n(x) = F(x)$), then we can make some general statements regarding the relationship between the long-run effective annual capability growth rate and the behavior of acquisition stakeholders. These are captured in the following theorems with the proofs provided in Appendix A.

First, let us consider the optimal technology policy when we assume that there is central control of the technology policy. We would like to know the behavior of the optimal symmetric technology policy with regard to the number of capabilities, n, provided by the system to be acquired.

Theorem 2.1 (Performance Decreases with Multi-mission Platforms)**.** *Given the above assumptions and assuming symmetric technologies, there exists a single optimal symmetric technology policy, $g^* \geq 0$, that decreases as the number of system capabilities, n, increases. Consequently, the effective capability growth rate declines as well.*

What this tells us is that there is an optimal technology policy that maximizes the growth in capability over time. However, there is a price to be paid for acquiring multi-mission platforms. The more technology development activities required, the greater the probability that one or more will delay the acquisition program. Consequently, all else being equal, the expected duration of the acquisition program increases. To compensate, the optimal target for each capability must be decreased versus a system that provides only a single capability. This is not to suggest that there are not other benefits to multi-mission or multi-capability platforms. It just means that some performance will have to be sacrificed to achieve the benefits of a multi-mission platform be it cost savings or otherwise.

Immediately following from Theorem 2.1 is a corollary regarding the impact of the integration time, I.

Corollary 2.1. *As I increases, g^* increases and the long-run effective annual growth rate decreases.*

Essentially, the presence of overhead in the form of system integration time diminishes the benefit of faster cycles. As this overhead increases, a greater leap in capability is required to compensate for the delay. This suggests that for very complex systems that require extensive integration time, it may, in fact, be preferable to set higher capability targets than for simpler systems.

Next, we reach the key finding. When there are independent stakeholders that have influence over the program requirements, a tragedy of the commons occurs. Defense acquisition programs, as with any public program, typically serve a number of stakeholders. Each of these stakeholders may have different objectives and, hence, different requirements for the acquisition program. The acquisition program serves as a common resource for stakeholders to achieve their individual objectives. Since there is essentially no cost to the stakeholder (the program is funded with public money), they are incentivized to demand more aggressive capability targets than is optimal.

34

This is stated more precisely in the following theorem.

Theorem 2.2 (Acquisition as a Tragedy of the Commons). *Assuming that the technologies are symmetric and each capability is supported by an independent group of stakeholders, the open loop equilibrium results in a technology policy that is more aggressive than the optimal policy (i.e., $g > g^*$), and consequently, the long-run effective annual capability growth rate is less than optimal.*

Since this is the key result of this chapter, it merits some additional discussion. First, consider the stakeholders. Who are they? The most obvious are the warfighters themselves, but they also may be defense contractors, members of Congress, or even participants in the acquisition programs themselves. The key is that program requirements are in some sense negotiable and open to influence as each party attempts to maximize its own objectives. For example, different groups of warfighters may push for more advanced capabilities that benefit themselves more than others. For instance, if we consider a multi-mission surface combatant, the Marines may desire an improved shore bombardment capability while the Navy may desire an improved fleet air defense capability. Or perhaps a Congressman may push for a particular advanced technology because it will mean additional long-term employment in his district. Regardless, the effect is the same.

We can study the behavior of non-cooperating parties using game theory. Theorem 2.2 essentially states that Nash equilibrium of the technology policy is not the same as the optimal technology policy. This means the optimal solution is unstable even assuming stakeholders agree to cooperate. We can imagine it this way. Let us assume that for a particular program all of the stakeholders agree to the optimal technology policy. But if everyone follows the optimal policy, then it is in the best interest of any one stakeholder to deviate from the policy. That stakeholder would be better off to push for a little more capability while everyone else follows the cooperative policy. Of course, if one deviates, then it is in the best interest of the others to deviate as

well. Very quickly we end up at the Nash equilibrium.

For example, let us assume that we are acquiring a surface combatant that provides capabilities for both anti-air and subsurface warfare. Those requiring the anti-air mission may push for additional capability, say a more advanced radar system. But opting for the more advanced radar system will likely mean a longer development time. This means that those requiring the subsurface warfare capability must wait longer than anticipated. If they must wait, they might as well demand a more advanced sonar. Of course, the interaction effect between the two technology development efforts means that the program will likely be even longer now causing both parties to desire even more advanced technology to compensate for the delay. Eventually, a stable point is reached where the cost of the delay outweighs the gain from additional capability.

This situation is quite familiar in defense acquisition. As the expected length of an acquisition program increases, more capability is demanded to compensate for the delay. In other words, the new system had better be worth waiting for. But really, the most important finding of this analysis arises when we consider Theorems 2.1 and 2.2 together. That is given the characteristics of the technologies involved, the complexity of system integration, and the structure of the program, there is an optimal technology policy, and it may not always be the modest policy that evolutionary acquisition would recommend. However, when there are multiple stakeholders with different objectives, the technology policy pursued will likely be more aggressive than optimal. Thus, we see worse performance over the long-run than we would otherwise expect. Better performance could be achieved if all stakeholders cooperated, and each sacrificed a small amount of capability to speed the completion of the program. This is unlikely, however, because there is always an incentive to deviate.

There is one final feature of the model that is of interest. The following theorem deals with the impact of additional capabilities on the gap between the cooperative

and competitive technology policies.

Theorem 2.3 (Multimission Platforms Exacerbate the Competition Effect). *As the number of capabilities with independent stakeholders, n, increases, the gap between the optimal and competitive policies increases.*

What we find is that as the number of capabilities, n, increases, the optimal solution is to sacrifice a little more capability to achieve better performance over time. However, the opposite is true for the competitive policy. When stakeholders do not cooperate, the technology policy becomes more and more aggressive because the interaction effect is exacerbated, and all participants are increasingly worse off.

2.4 Numerical Example

To illustrate the implications of the mathematical model more concretely, a notional example will be presented. First, we will assume that the expected duration of an individual technology development activity, $W(g)$, is governed by an exponential function. In particular, we assume

$$W(g) = e^{2g} - 1.$$

This function is depicted in Figure 2.4 as the $n = 1$ curve. We will also assume that the duration of each individual technology development activity is exponentially distributed. The exponential distribution is convenient because the distribution of the maximum of n independent and identically distributed random variables can be found analytically. However, the analysis would work just as well with another non-negative distribution such as lognormal or beta. Figure 2.4 depicts the expected duration of the technology development phase versus the number of capabilities provided by the system, n.

We see that curves increase with n but appear to be converging to a limiting case. This is, in fact, what is happening since extreme value theory tells us the distribution

37

Expected Duration of Technology Development Phase
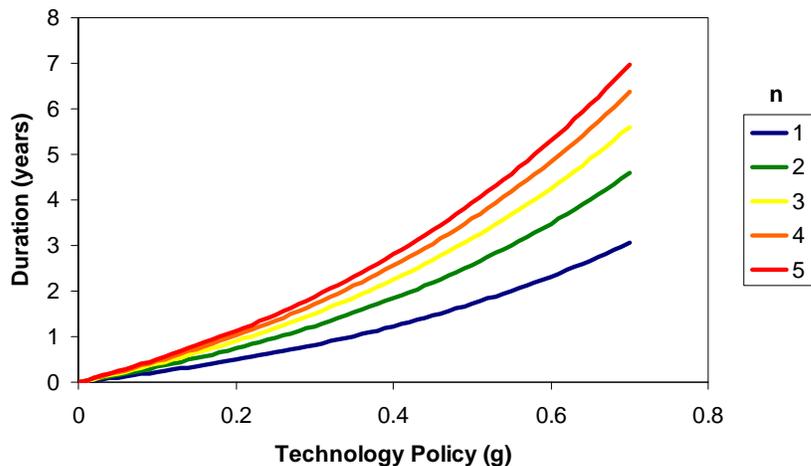
Figure 2.4: Expected duration of the technology development phase versus the number of capabilities, n.

of the maximum of n normalized i.i.d. random variables converges to one of three distribution types as $n \to \infty$. So we see that as n increases we would, all else being equal, expect to see the duration of our acquisition programs increase.

Given this structure for the expected duration of the technology development phase, what is the outcome for any given technology policy, g? To determine the relationship between the technology policy and n, we must specify an integration time. In this example we will assume that integration time, I, is three years. When we do so we obtain Figure 2.5.

Figure 2.5 reveals that the long-term effective growth rate is unimodal as was implied by Theorem 2.1. Thus, there is a single optimal policy, g^*, for each value of n. We also see that the effective growth rate decreases as n increases and that mode shifts to the left. This is also consistent with Theorem 2.1.

Now we would like to compare the optimal policy to the competitive policy (i.e., the Nash equilibrium). This comparison is presented in Figure 2.6. Here we see that when $n > 1$ there is a gap between the optimal policy and the competitive policy,

Effective Growth Rates
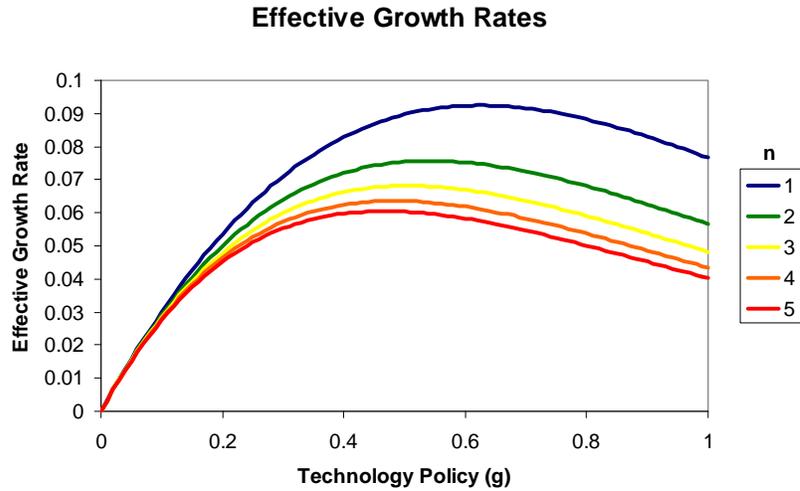
Figure 2.5: The long-term effective annual capability growth rate versus the number of capabilities, n.

hence the tragedy of the commons. We also see that as n increases the optimal policy decreases slightly, but the competitive policy increases rapidly. Thus, we see the gap widens as n increases. This example clearly illustrates how the technology policy becomes more aggressive when stakeholders do not cooperate.

Technology Policy vs Number of Capabilities
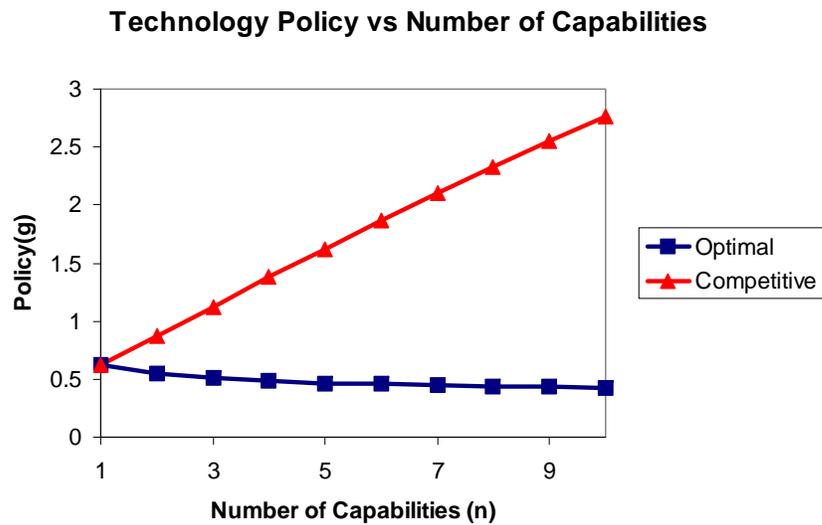
Figure 2.6: The optimal technology policy, g^*, and the competitive technology policy versus the number of capabilities, n.

We can also consider the corresponding long-run effective capability growth rates. These are depicted in Figure 2.7. In this figure we see that capability growth for the optimal policy declines modestly as n increases, but that it rapidly approaches zero for the competitive policy. Thus, we see that the penalty for a lack of stakeholder cooperation can be quite severe.

Effective Growth Rate vs Number of Capabilities
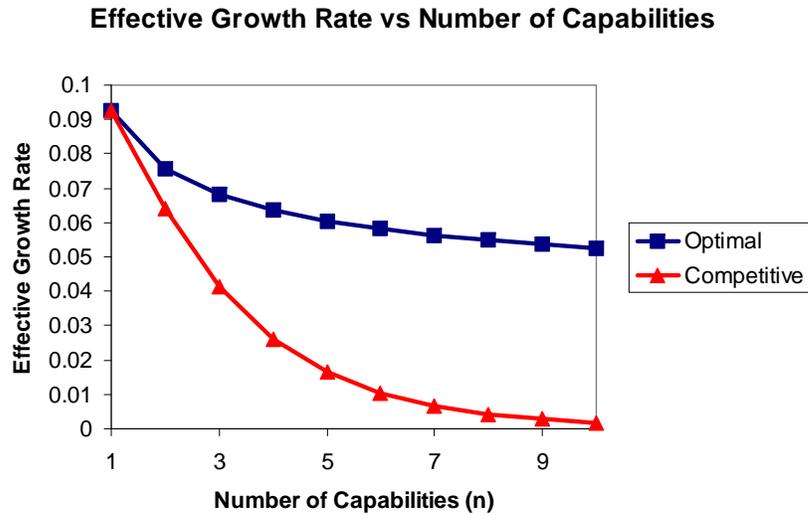
Figure 2.7: The optimal effective growth rate and the competitive effective growth rate versus the number of capabilities, n.

What is the impact of the integration time? To find out, we fix the number of capabilities at $n = 3$, and then vary I. Figure 2.8 displays the relationship between the technology policy and I. We see that as I increases, the technology policy for both the optimal and competitive cases increases, but at a diminishing rate. More interesting is the impact on the effective growth rate. If we examine Figure 2.9, we see that as I increases the effective growth rate also decreases for both, but much faster for the optimal case. In fact, the two appear to be converging as the impact of the integration time begins to dominate.

The above examples provide us some sense of the behavior of acquisition programs with respect to the number of capabilities provided and the time required for

Policy Sensitivity to Integration Time
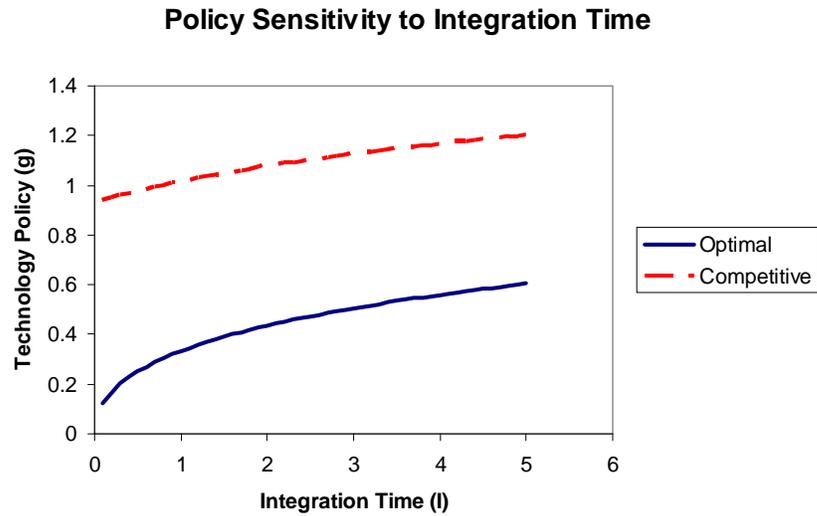
Figure 2.8: Sensitivity of the technology policy to the integration time.

Growth Rate Sensitivity to Integration Time
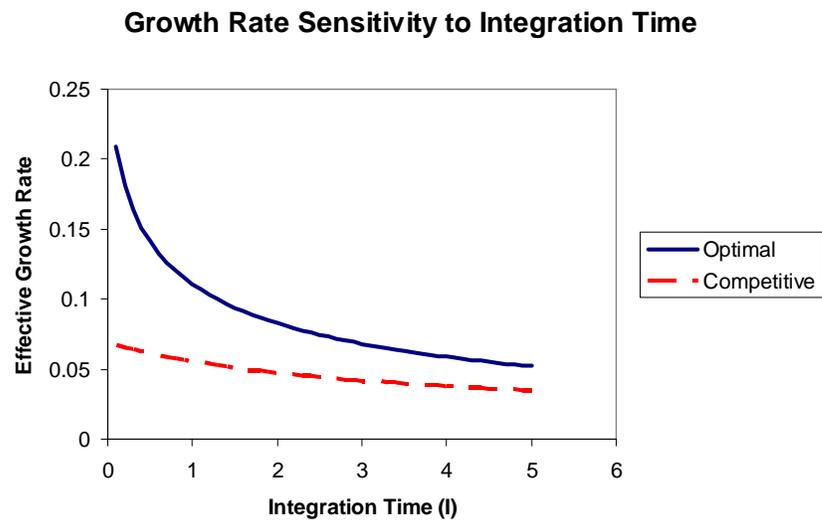
Figure 2.9: Sensitivity of the effective annual growth rate to the integration time.

integration. More importantly, however, it demonstrates the impact of competitive behavior on the performance of the acquisition system. Of course, all of the above results focus on the long-term performance. What might also be of interest to policy makers is the impact of competitive behavior over the short-term and the level of risk associated with program performance.

To capture short-term behavior we require another metric. We will consider the level of deployed capability, $C(t)$, at a given time, t. Since this is a stochastic quantity we will consider both its expected value and its variance. Unfortunately, $C(t)$ is difficult to describe analytically even when the underlying activity distributions are exponential, but Monte Carlo simulation can be used to estimate both $E[C(t)]$ and $\text{Var}(C(t))$.

The first question we would like to consider is how different technology policies affect the uncertainty in the level of capability deployed. Figure 2.10 depicts the expected level of capability deployed in year 20 when the integration time is three years, $I = 3$, and the system provides three capabilities, $n = 3$. The dashed lines constitute one standard deviation above and below the expected value. Note that the level of deployed capability is unimodal with respect to g and that the peak corresponds with the optimal policy found in the long-run analysis above. Also, note that the the level of uncertainty increases with g, though at a diminishing rate. In other words, as the technology policy becomes more aggressive, the uncertainty in the level of capability actually delivered increases.

If we take a closer look at the standard deviation in Figure 2.11, we see that, actually, the standard deviation does not strictly increase. In fact, it actually decreases after a certain point. This is quite reasonable because as g increases eventually it becomes unlikely that any programs will be completed within 20 years. Thus, the uncertainty must decrease. It is important to note, however, that if a decision maker

42

Capability vs Technology Policy

n = 3, *I* = 3, *t* = 20
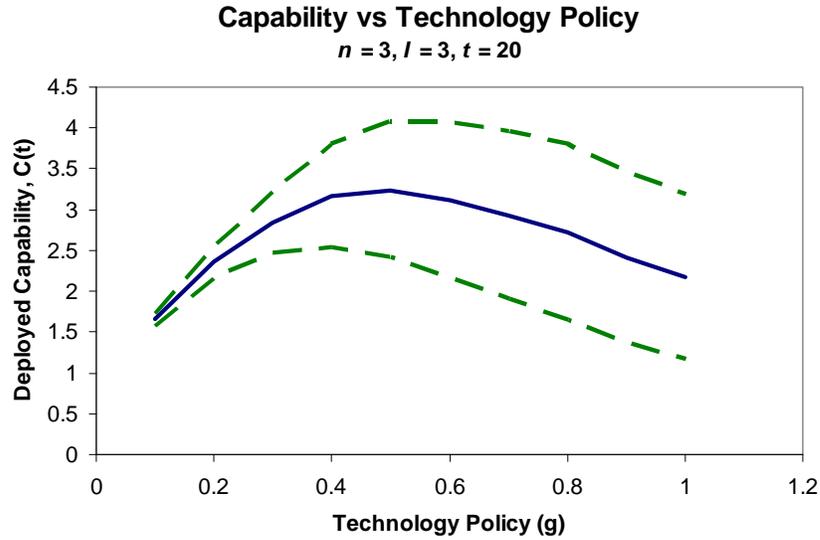
Figure 2.10: The level of deployed capability after 20 years when $I = 3$ and $n = 3$. The solid line is the expected value and the dashed lines are one standard deviation from the mean.

wanted to trade expected capability for reduced risk, he would likely choose a technology policy that is less aggressive than the optimal policy. Most of the policies above the optimum are dominated, and thus, it would appear that the competitive policy is also dominated.

Standard Deviation vs Technology Policy

n = 3, *I* = 3, *t* = 20
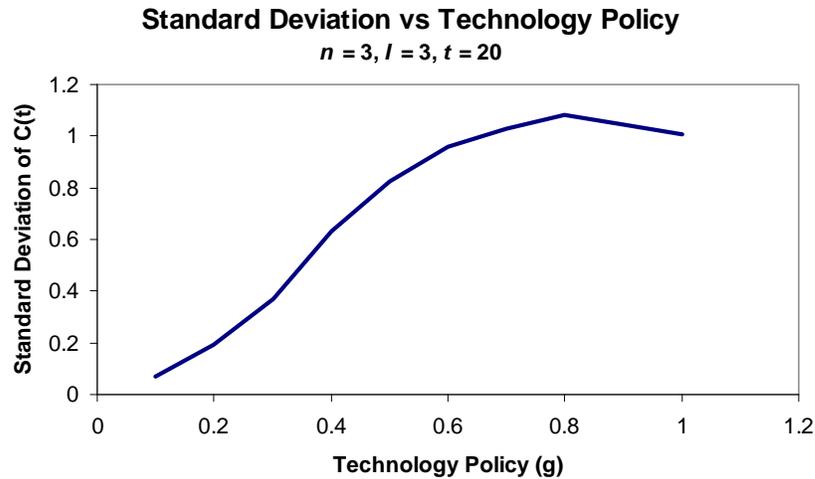
Figure 2.11: The standard deviation of deployed capability after 20 years when $I = 3$ and $n = 3$ versus the selected technology policy, g.

One feature of the model worthy of mention is that uncertainty actually decreases as the number of capabilities, n, increases. This is illustrated in Figure 2.12 and is not unexpected. This is a well known result in extreme value theory. As $n \to \infty$, the probability that at least one technology development activity takes the maximum possible amount of time approaches one. Thus, as the number of system capabilities increases, uncertainty decreases in that the acquisition program will undoubtedly take a very long time.
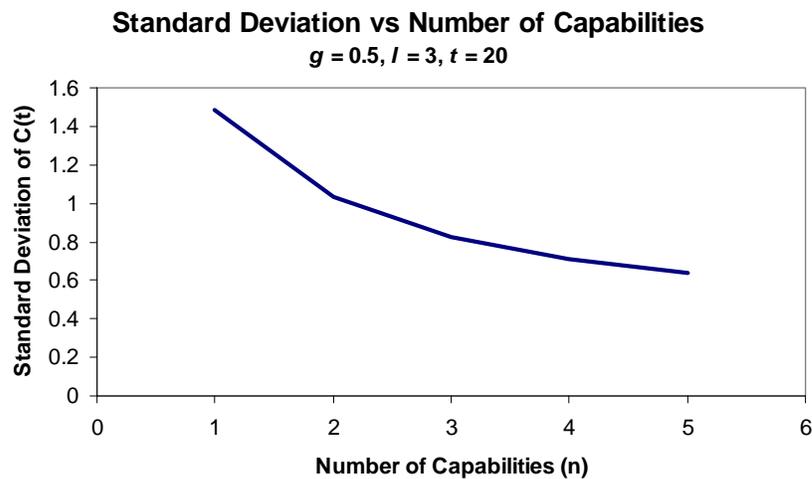
Figure 2.12: The standard deviation of deployed capability after 20 years when $I = 3$ and $g = 0.5$ versus the number of system capabilities, n.

The above results would seem to suggest that even in the transient case, the long-run optimal policy outperforms the competitive policy. This is confirmed by Figure 2.13 which shows that from beginning to end, the optimal policy outperforms the competitive policy in terms of the expected deployed level of capability.

We can link this short-term analysis to the previous long-term analysis by calculating the effective annual capability growth rate versus time. We find that when we do so, as t increases, its expectation approaches the value predicted in the long-run analysis quite quickly (Figure 2.14). In fact, the standard deviation of the effective annual capability growth rate decreases with time, so it appears to be quite reasonable
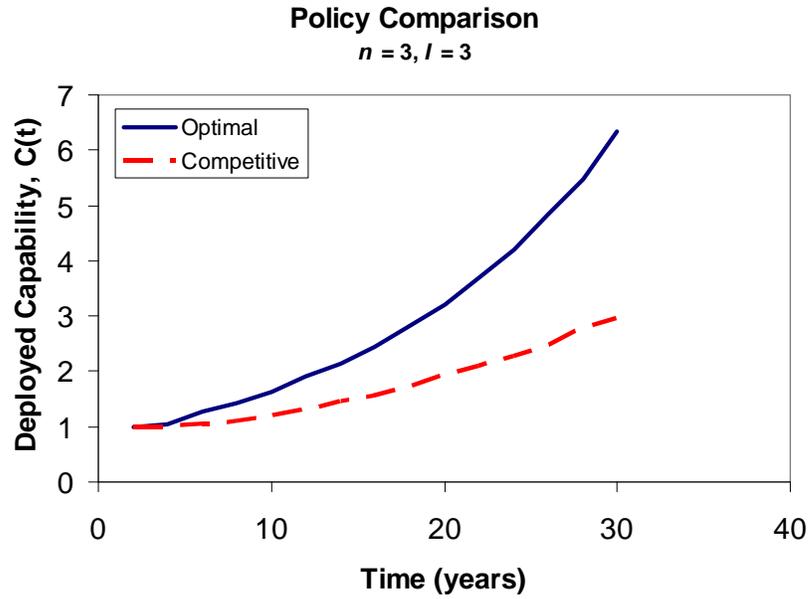
Figure 2.13: The expected deployed capability versus time when $I = 3$, $n = 3$, and $g = 0.5$.

to use the long-range results as a basis for analyzing policy implications.
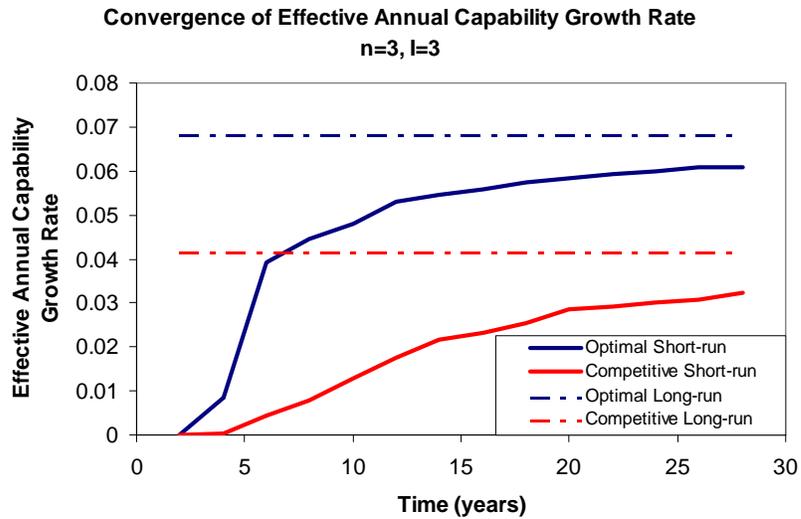
Figure 2.14: A comparison of the effective annual capability growth rate between the short-run and the long-run when $I = 3$ and $n = 3$.

This example has revealed that when stakeholders in the defense acquisition system pursue independent agendas, there can be a potentially substantial negative impact on the level of capability actually deployed to the field. Of course, this analysis is fairly abstract and based on assumed parameter values and probability distributions, and thus, we should not take the numbers literally. However, Section 2.3 revealed that the results are fairly general and do not depend on the particular distribution assumed. Consequently, what we can draw from this example is that the structure of acquisition programs incentivizes participants in the acquisition system to behave in a manner that causes the system to underperform, and this problem is exacerbated as the number of capabilities provided by a system increases.

2.5   Policy Implications

What we can conclude from this analysis is that, from a performance standpoint, every acquisition program has some optimal technology policy that is dependent upon the nature of the system and technologies involved. Unfortunately, the implementation of this optimal acquisition strategy is not trivial. The DoD's increased emphasis on multi-mission or multi-capability platforms may lead to operational cost savings and increased flexibility, but it creates a tension between the competing missions and capabilities. A multi-mission platform means that some capability must be sacrificed relative to a specialized system in order to deliver the system in a reasonable time frame and maintain the optimal acquisition strategy. The result is that the optimal strategy requires an unstable technology policy that incentivizes stakeholders to deviate from that policy. Thus, there is a tendency in the Department of Defense to pursue an overly aggressive technology policy.

The multi-mission problem has certainly been noted before. In fact, a recent RAND study analyzing cost growth in US Navy ships suggested acquiring smaller

mission-focused ships over large multi-mission ships as a strategy for restraining requirements growth [1]. Instead, the contribution of this work is understanding why a multi-mission system leads to excessive capability goals, and understanding this "why" is important to creating a rational acquisition policy. The problem is not intrinsic to multi-mission platforms, but rather, sub-optimal performance is a result of the rational behavior of participants within the acquisition system. Evolutionary acquisition polices were instituted to address this problem. However, compliance is largely voluntary.

In as much as the optimal policy tends to be more moderate than the stable policy, we can say that the former is more evolutionary while the latter is more revolutionary. The implication is that while evolutionary acquisition is more appealing from a performance standpoint, revolutionary acquisition is the more natural outcome. This means that the Department of Defense cannot expect programs to voluntarily comply with evolutionary acquisition procedures since the nature of the system pressures programs towards revolutionary leaps in technology. Consequently, if the DoD is serious about evolutionary acquisition, technology maturity requirements must be strictly enforced.

More broadly, however, the results presented in this chapter support the assertion, at least regarding evolutionary policy, that the participants in the defense acquisition system are disincentivized to implement or follow acquisition reforms. In the example presented here, the DoD established a process to comply with reforms suggested by the GAO and others, and then proceeded to ignore them in practice despite agreeing in principle. The tragedy of the commons presented here explains, in part, why this apparently contradictory behavior occurs, and opens the door to remediation of such phenomena.

Unfortunately, countering such a tragedy of commons is non-trivial in a government context. While private firms may overcome similar situations through the use

47

of monetary incentives, such a solution seems less plausible for the defense acquisition system. It is unlikely that monetary incentives would be legal or even effective. Often stakeholders in defense have strong beliefs regarding the import of desired systems and capabilities [41]. They are not likely to be swayed by material compensation. It may be that strict regulation and oversight is the only solution to countering the tendency to defy acquisition reform.

CHAPTER III

A SYSTEMATIC ANALYSIS OF THE COST AND PERFORMANCE IMPACT OF ACQUISITION TECHNOLOGY POLICY

Over the past several years, the United States Department of Defense (DoD) has been attempting to transform itself from an organization designed to meet the Cold War threat of the Soviet Union to a more flexible, adaptable organization that is ready to meet the regional and asymmetric threats the US expects to face in the coming years. To facilitate this transformation, several modifications have been made to the defense acquisition system, the most important being the shift to evolutionary acquisition.

Evolutionary acquisition is an attempt to address one of the most common criticisms of the defense acquisition system. As has been discussed in previous chapters, traditional acquisition programs attempt large, revolutionary leaps in system capability through the use of immature and risky technology. Not only does immature technology often require more time and money to develop, but it also introduces uncertainty that may lead to significant delays and cost overruns. Consequently, warfighters must often make due with increasingly obsolete equipment during the long intervals between new system deployments, and there is little flexibility to adapt to emerging threats and exploit technology opportunities.

Evolutionary acquisition, on the other hand, attempts to set more modest capability goals for each acquisition. The idea is to use more mature, and hence, less risky technology, in order shorten acquisition cycle times. Thus, each acquisition cycle under evolutionary acquisition should be shorter and cost less that more traditional programs. As a result, warfighters should receive more frequent upgrades to their

equipment and, thus, should be at less risk of going to war with obsolete hardware.

Despite the apparent motivation to implement evolutionary acquisition and committing the approach to policy, it would seem that the DoD has had limited success in doing so [47]. In fact, the US Government Accountability Office (GAO) has suggested that DoD reforms have not gone far enough [30, 31, 33, 36]. They advocate adapting commercial best practices regarding technology and product development to the defense acquisition system. Among these are a centralized portfolio approach to managing new systems and technologies, a staged knowledge-based approach to both acquisition and technology development, strict enforcement of technology maturity requirements, and a more evolutionary approach to new system development.

Since these reforms are derived from the commercial world, the obvious question is whether they will translate well to a government context. The defense acquisition system differs from a commercial product development process in several respects. In particular, the government essentially serves as a technology developer, system developer, customer, and user. Furthermore, the DoD along with a few allies are really the only customers for the systems and technologies developed within the defense acquisition system. Thus, there is a more limited capacity to purchase multiple evolutionary iterations of a system than there would be with a consumer product. Consequently, the pertinent question is, if evolutionary acquisition were fully implemented, would there be any tangible benefit for the Department of Defense? As was asserted in Chapter 1, it appears that policy makers at the DoD and elsewhere often do not fully understand the trade-offs inherent in altering a system as complex as the defense acquisition system.

As a first step to better understanding the system level trade-offs of technology policy on acquisition, the work presented in this chapter attempts to model the basic "physics" of the acquisition system, in particular the relationship between the R&D process and the acquisition life-cycle. The purpose is to gain insight into the most

fundamental system-level influences on the efficacy of acquisition policies. To that end, an idealized view of the acquisition system is adopted to which complicating factors may be subsequently added to test their effects. The acquisition model was implemented as a discrete event simulation with the key decision variable being the maturity level at which a technology moves from R&D to an acquisition effort. Extensive sensitivity analyses were performed and several insights into the impact of technology policy on acquisition were generated. The most important output of this effort, however, is an informed set of future research directions that will facilitate more definitive answers to major policy questions regarding evolutionary acquisition.

3.1   Background

As was mentioned previously, evolutionary acquisition is an attempt to reduce acquisition cycle times by setting capability goals that are more modest than is typical of a traditional program. This allows programs to utilize more mature technology and, hence, reduce the amount of technology risk. In theory, this should reduce cost, schedule, and performance uncertainty. The hope is that it will lead to less expensive acquisition programs that proceed more quickly. Consequently, warfighters would receive up to date equipment more frequently and at lower cost.

Evolutionary acquisition was instituted at the US Department of Defense in 2003 following a revision of DOD Directive 5000.1 and DOD Instruction 5000.2 [18, 19]. In particular, the Defense Acquisition Guidebook indicates the evolutionary acquisition is the preferred acquisition approach [20]. According to DODI 5000.2, Section 3.3 [19], there are two ways to implement evolutionary acquisition, incremental and spiral development. Essentially, one large acquisition program is created, but broken into smaller pieces. Each of these pieces effectively functions as a small acquisition program. Under incremental development, the sub-programs are pre-planned. There is a target set of capabilities that the final system should provide, but it is achieved

through a series of phased deployments, each more capable than the last. Under spiral development, however, the phased deployments are undefined. Instead, as each spiral is deployed, feedback is collected from the users and used to shape the capability goals of the next spiral.

Since spiral development is nominally the preferred approach for implementing evolutionary acquisition, several studies have been conducted on its efficacy. The defining characteristic of spiral development in defense acquisition is concurrency. The spirals overlap, yet are interdependent. The result, according to Dillard and Ford, is that the first delivery of capability is achieved more quickly than under a traditional program, but the overlapping spirals may lead to a substantial amount of rework and backlogs that results in a slower delivery of objective capabilities [16]. Furthermore, a RAND study of the implementation of evolutionary acquisition programs found that spiral development introduces significant management difficulties and that the user feedback process often resulted in confused and contradictory requirements for future spirals [47]. As a result, many programs that started out using spiral development ended up reverting to incremental development.

Regardless of the approach taken, the motivating issue behind evolutionary acquisition is cycle time. In theory, shorter cycles mean that each is less expense and new technologies can be moved into the field faster to meet emerging warfighter needs. The driving issue, then, is really how big of a leap in capability should one attempt during each acquisition cycle? Of course, the risk associated with the size of the leap is linked to the maturity of the required technology. Thus, evolutionary acquisition is really all about technology policy because with a large enough leap, evolutionary becomes revolutionary.

So where does the DoD's approach to evolutionary acquisition come in? A key issue is that the DoD does not manage technology or "product" portfolios in same manner as a large commercial enterprise. In part, this is due the public nature of the

defense enterprise. Even so, the GAO asserts that the DoD should adopt additional commercial best practices regarding the centralized management of its acquisition and technology portfolios and the management of technology transitions from R&D to acquisition programs [31, 36]. Under the current system, there is often a funding gap in technology development. Early stage technologies are funded through the R&D system (or S&T as it is known in DoD) and late stage technologies are often funded in support of a particular acquisition effort. It is technologies in the middle stages of maturation that are often left without obvious ownership and hence funding. Consequently, if certain technologies are required by an acquisition effort, their development through the middle stages must be funded in support of the development of a particular system. This requires early commitment to a technology when its final realization is still uncertain. In the past, this has often led to disappointment as technologies took longer to develop and did not perform as well as expected. Theoretically, if the DoD adopted the commercial new product model that the GAO suggests [31, 36], it would allow the DoD greater flexibility in how to select and mature technologies for development in anticipation of future acquisition program needs. This would reduce the burden and risks of technology development on acquisition programs since they could choose from a portfolio of mature technologies.

So in the end, the two fundamental questions of evolutionary acquisition are how mature should technologies be when they are transitioned from R&D to acquisition efforts, and what is the best approach to mature them? All else being equal, this essentially determines the acquisition cycle time as well as the size of the capability improvement for each cycle. Ultimately, the answer will hinge on factors such as the cost of technology maturation, the rate of learning from fielded systems, and the overhead cost associated with an acquisition cycle.

3.2   Model Setup

The motivation behind the structure of the model is to represent the set of commercial best practices recommended by the GAO for implementation in the context of the defense acquisition system. This includes both a staged, centrally managed technology development process as well a strictly enforced acquisition program life-cycle. Given the staged nature of both R&D and acquisition, discrete event simulation was the logical choice to capture the behavior of the system. As was mentioned previously, the representation of the defense acquisition system presented here is intentionally scaled-down and idealized. The benefit of an idealized model is two-fold. First, the scaled-down representation is more tractable and allows us to attempt multiple experimental excursions. Second, it allows us to consider the structural impacts of technology policy unobscured by the inconsistent implementation that occurs in the actual defense acquisition system. In particular, the modeling emphasis was on the linkage between the movement of technologies through the R&D process to the length and cost of the acquisition cycles. In order to represent the impact of technology policy on defense acquisition, there were three key features of the system that required consideration: the movement of technologies through the R&D system, the movement of programs through the acquisition process, and the rate of technological progression.

The simulation was implemented using the Arena 10.0 software package and consists of three major components: the technology development process model, the system acquisition process model, and the technical progress model. The technology development process model describes how technologies with potential defense application are matured through the defense R&D system. This process provides a portfolio of technologies for use by acquisition programs. The system acquisition process model describes the life-cycle of a defense acquisition program from concept development to deployment. Finally, the technical progress model describes how the capabilities provided by technologies improve over time.

3.2.1 Technology Development Process Model

The technology development process model simulates the movement of individual technologies through a maturation process. Ideally, a technology development process is centrally managed and staged. Technologies are selected for development based on their potential applicability to future products. In the commercial world, product and technology roadmaps drive development. These roadmaps, and the organization's commitment to them, provide a shared vision that DoD often lacks. However, developing the technologies to satisfy the roadmap entails a certain amount of risk. In order to mitigate that risk, each technology must pass through a series of stage-gates. Each gate provides an opportunity to evaluate the status of a technology and determine whether or not it should continue to receive funding. Such a system facilitates prioritization of technology projects as well as risk mitigation. It is important to note that the Department of Defense has not consistently implemented such a system [31]. Instead, there are a number of different organizations throughout the DoD that perform or fund R&D work, each with its own way of managing technology projects. These inconsistencies preclude the effective management of technology development and promote duplication and mismatch between the technology supplied by R&D organizations and the technology demanded by acquisition programs. Consequently, for this study, the technology development process was modeled in the spirit of the GAO's recommendation of a centrally managed and staged technology development process.

The process starts when new but immature technologies arrive for evaluation. The arriving technologies are prioritized and then funded until the budget is expended. Technologies that are rejected are considered for funding in future rounds, and those that are successfully matured move on to the next stage. The sequence repeats until each technology is either successfully matured or discarded. The maturity of each technology is measured by the Technology Readiness Level (TRL) scale.

The TRL system is a qualitative scale for measuring technology maturity that was developed by NASA. Technologies at the lowest level of maturity receive a score of 1 while those at the highest level receive a score of 9. It has been recently adopted by the DoD as the standard measure of technological maturity in the defense acquisition process. There are several known issues with the TRL scale including inconsistent application, its inability to account for integration risk, and a hardware rather than software focus [14,55,69,71]. Even so, the TRL scale does provide a convenient means to roughly categorize the maturity of a technology.

Besides a maturity level, the technologies in the model have a few other attributes relating to the cost to mature them, their expected contribution to system capability, and the area of application. Since these attributes are assigned randomly, technologies arriving at the beginning of the technology development process will vary considerably in their costs, application, and performance. The purpose of a properly functioning technology development process is to prioritize and fund these technologies by potential cost and benefit. The process used in this simulation is represented in Figure 3.1.
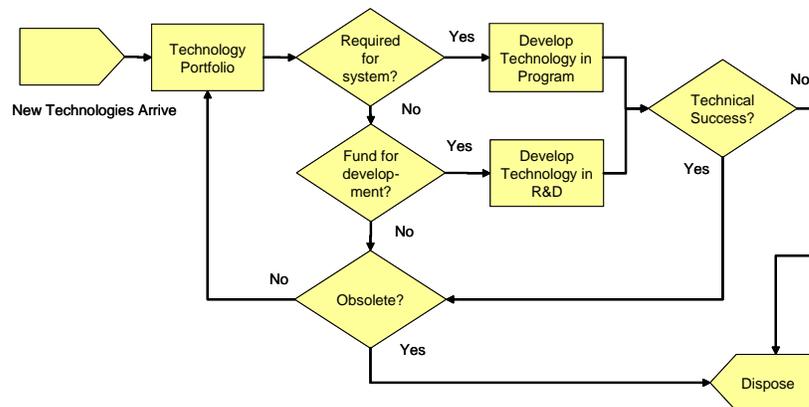
Figure 3.1: The Technology Development Process Model

The steps in the process are as follows:

1. Technologies arrive randomly from exogenous sources and are assigned attribute values.

2. Technologies are collected in the technology portfolio which constitutes the set of technologies available for use in military systems.

3. On an annual cycle, each technology is considered for maturation. If a technology is required by an acquisition program, and it meets minimum maturity requirements, it is sent to the acquisition program for further development. Otherwise it is considered for R&D funding (or S&T funding as it is known in the DoD).

4. The technologies are sorted by TRL level, and there is a separate budget for maturing each level. Funding technologies for maturation at each level is a knapsack problem. Instead of solving the computationally intensive knapsack problem, a well-known heuristic is used. The technologies are sorted by benefit/cost ratio and funded in order until the budget is depleted. In this case, the benefit is the technology's performance level, and the cost is the expected development cost.

5. If a technology is not funded, it is evaluated for obsolescence. Obsolescence is defined as having a lower performance level than the best technology currently deployed in the same application area. If the technology is not obsolete, then it is returned to the portfolio for future consideration.

6. If a technology is funded for development, it encounters a variable delay based on its TRL level. Cost is accumulated at a rate determined by the TRL level multiplied by the technology's base cost.

7. After the development delay, it is randomly determined whether or not the technology development effort succeeded. The probability of success is determined

57

by the TRL level.

8. If the development effort fails, the technology is dropped.

9. If it succeeds, the TRL level is increased by one, and it is checked for obsolescence and returned to the portfolio.

It is also important to note the following assumptions regarding the technology development process:

- Technologies assigned to a program for development must go through each development stage in series, but it is assumed that the technology is pre-approved for funding for all stages.

- All technologies to be used in a system must reach TRL 7. If technology development fails at any stage, the technology is dropped, and the program must find a replacement. If development is successful, then the technology is returned to the portfolio at TRL 7.

- TRL 7 is the highest achievable level in this model, and no further development is required for these technologies (This is because TRL levels beyond 7 are really system specific [45]).

- It is assumed that a technology either achieves its predicted performance or fails entirely. This assumption is primarily for model tractability while still capturing the uncertainty inherent in technology development.

- There is no budget discipline for technology development. Technology development efforts are allowed to under run or overrun their budgets. Thus, a development effort that finishes under budget will generously return unused funds, and an effort that overruns is provided the resources it needs to reach a conclusion. Budget discipline in the R&D system is avoided because it allows

for direct comparison with the cost of technology development in acquisition programs where overruns are often allowed. Furthermore, it allows us to focus on the true impact of evolutionary acquisition policies without the confounding effects of budgetary politics.

3.2.2 System Acquisition Process Model

The system acquisition process model describes the life-cycle of a defense acquisition program. The nominal five stage process is depicted in Figure 1.1. There are also three key decision points in the process called milestones. Technically speaking an acquisition program does not begin until after milestone B. But in practice, "proposed" programs develop a great deal of momentum after they are approved by the Joint Requirements Oversight Council (JROC), and the acquiring service begins programming and budgeting activities immediately [36, 41]. Furthermore, milestone requirements are often bypassed so even though a program nominally moves through all phases sequentially, in practice, there is often a great deal of concurrency between the technology development, system development, and production phases. As with the technology development process model described above, the acquisition process model in the simulation will assume that acquisition programs follow the rules, and consequently, programs will move through each phase in order with no concurrency.

Within the simulation model, the basic unit in the system acquisition process is a program to acquire a system. It is assumed that DoD has several different types of systems. Each type is continuously cycling through the acquisition process. For example, if the Air Force deploys a new air superiority fighter, it is assumed that it will begin concept development of its replacement shortly after. This assumption will be relaxed later.

Each type of system is dependent upon several technologies, each from a different application area. For example, an air superiority fighter might require a propulsion

technology, a sensor technology, and an avionics technology. The acquisition process model used in the simulation is illustrated in Figure 3.2. The steps in that process are as follows:
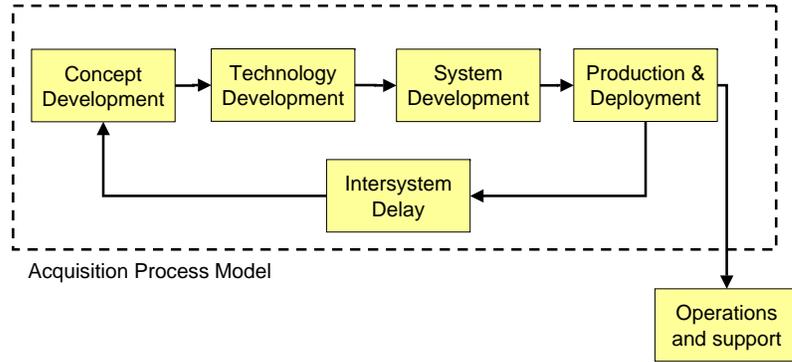
Figure 3.2: The System Acquisition Process Model

1. When a system enters concept development, it encounters a randomly determined delay and accumulates cost at a pre-specified rate.

2. Following concept development, the system enters technology development and requests technologies for each of its required application areas. The technology selection rule is as follows: a program selects the technology from a required application area with the highest performance that meets the minimum TRL requirement. The minimum TRL level is a simulation-wide parameter and applies to all systems.

3. After technologies are selected, their development proceeds in the manner described earlier (Section 3.2.1). If a technology fails, a replacement must be selected. The same selection rule applies as before, only now the minimum TRL level is the Fallback TRL level. The Fallback TRL level allows a program to select a more mature technology in the event that the desired technology failed.

4. The system is held in technology development until it has technologies at TRL 7 for all required application areas.

5. Once a system is released from technology development, it enters system development and, after that, production.

6. After production, the system is considered to be deployed. The system is held for a delay period (no cost incurred) and then returned to concept development to start the next acquisition cycle.

Some additional features of the of system acquisition process model are:

- For each phase there is a randomly determined delay and associated cost accumulation.

- If multiple systems share a technology requirement and enter technology development at the same time, they will share the technology development effort.

- The Operations and Support phase is not costed in the model since it is outside of the acquisition system in the context of this analysis.

Ultimately, the purpose of acquiring a system is to provide military capabilities. It is assumed that each system deployed provides a capability. Capability in the model is an abstract representation of military utility. It is assumed that there is a synergistic effect between the technologies employed in the system. That is the system is greater than the sum of its parts. Thus, a multiplicative model is used to represent capability. The capability of a deployed system is the product of the performance levels for each of its required technologies. Thus, an air superiority fighter without a propulsion system is useless no matter how capable its sensor is. This measure of capability allows us to determine the cost effectiveness of a particular technology policy.

3.2.3 Technical Progress Model

The final key feature of the simulation is the model of technical progress. Where do new, more capable technologies come from? It is important to note that the technology development model in this simulation does not consider basic research. In fact, TRL 1 signifies the transition of ideas, concepts, and technologies from basic research to applied research. Thus, we can assume that there is a certain amount of research occurring exogenous to the simulation. The source of this research may be from government or commercial sources. The key is that there is a constant inflow of new technologies and that their performance improves over time. The purpose of the technology development process is to adapt these technologies for use in a military system. There is one caveat, however, and that is that a purely exogenous technology progress model neglects the learning that inevitably occurs from fielding systems. For example, valuable information gathered from field use of a jet engine will likely inform the development of the next generation jet engine. Thus, there is a learning effect, and the more rapidly systems are fielded, the sooner subsequent learning will be available for future technologies. This is especially true for military specific technologies where the only source of user feedback is the military itself.

Consequently, the technology growth model in this simulation attempts to model both of these features. To do so, a hybrid model was created. First, there is a baseline technology coefficient for each application area. Whenever a technology is fielded, the coefficient is multiplied by a learning factor (e.g., 1.1). This captures the learning from implementation. Second, there is an exponential growth model for each application area. This represents the learning from exogenous R&D activities. The two are multiplied together to determine the current technology level and are represented by the equation

$$Ce^{gt}$$

where C is the technology coefficient and g is the exogenous growth rate. Arriving

technologies are assigned a performance as a random variation on this value. The parameters of this model can be adjusted to accommodate the specific situation of each application area. For example, technologies that are used commercially may have a high exogenous growth rate and low learning factor because their progress would continue regardless of military use. The reverse may be true for military specific technologies since there would be little learning from commercial use.

3.3 Experimental Design

3.3.1 Simulation Parameters

As previously mentioned, the DoD has been relatively inconsistent in its implementation of its own policies, and evolutionary policies, in particular, are fairly new. Consequently, using historical data to set simulation parameters is particularly problematic. In fact, a RAND study to assess cost growth in weapon system programs found a number of issues in the available cost data for defense acquisition programs [2]. Some of these issues include significant aggregation of data, baseline changes, changes in reporting guidelines, and incomplete data. The situation is worse for technology maturation. As indicated by the GAO, the DoD does not systematically track its technology development efforts [31]. Furthermore, the introduction of TRL levels to the DoD is fairly recent so there is little experience with their application in a DoD context. Since NASA has been using the TRL scale for some time, it would seemingly be a logical source of information regarding the cost and risk associated with maturing technologies through TRL levels. Unfortunately, a 2005 study at NASA to determine just that found that poor record keeping resulted in insufficient useful data to achieve statistically significant results [45].

Fortunately, the aim of this study is not to precisely recreate the defense acquisition system as it is, but instead to identify policy directions to determine how it should

be. This in combination with extensive sensitivity analysis allows for a more reasonable margin of error in setting the simulation parameters. Consequently, the actual values used in the experiments are an amalgamation from several sources including reports and studies from both government and commercial sources [6,21,22,26,37,45,72]. The baseline set of simulation parameters can be found in Appendix B, and a first order sensitivity analysis can be found in Appendix C.

3.3.2 Basic Experiment

In order to answer the research question posed in this chapter, three cases were developed. The three cases are variations on the key experimental variables, the Min TRL and the Fallback TRL. As mentioned previously, the Min TRL is the minimum maturity requirement for a technology used in an acquisition program, and the Fallback TRL is the minimum maturity selected when the first choice technology fails. The cases are as follows:

Base Case – The base case most closely resembles the current modus operandi of the defense acquisition system. Technologies are selected at mid TRL levels and final maturation occurs during the technology development phase of an acquisition program. High performing, but immature technologies are preferred over more mature, proven technologies. If a technology fails, however, the program will fall back to a more mature technology.

- Min TRL = 4
- Fallback TRL = 7

Evolutionary Acquisition – In this case, programs may only use fully mature technology. Maturation of technology is funded in the R&D system, and there is effectively no technology development phase.

- Min TRL = 7

Revolutionary Acquisition – Programs target maximum performance at all costs and, thus, always choose the most promising technologies. When a technology fails, another top performer is selected in its place.

- Min TRL = 4

- Fallback TRL = 4

There are several outputs of interest. These are the cost of operating the entire acquisition system, the cost of an individual program, the annual capability growth rate, and the acquisition program length. Of course, we are interested in the long-run behavior of these outputs. Consequently, to perform the experiments, the simulation was run for a warm-up period in order to fully populate the technology portfolio, and then statistics were collected on the outputs of interest.

In particular, each simulation was run for a warm-up period of 50 years and then statistics were collected for another 150 years. There are 40 replications for each experimental case. As for the acquisition programs, there are three system types each requiring three technologies. Each of those technologies falls into one of six application areas. It was assumed that the three acquisition programs are homogenous in terms of cost and schedule risk, and it was also assumed that the application areas are all homogeneous in terms of cost, schedule, and technical risk. The budget for the technology development process was set to $3 billion, and was allocated among the six stages so as to ensure a smooth flow of technologies through the system. It was also assumed that all of the stages are of equal length. This is simply to focus on the technical risk for the basic experiment. Finally, the technical progress model is identical for all six application areas and features a mix of exogenous technical progression and learning.

3.3.3 Sensitivity Analysis

The simulation developed is quite flexible and many different scenarios can be analyzed. As mentioned previously, a first order sensitivity analysis was performed, and the results are presented in Appendix C. It was found that the simulation outputs were particulary sensitive to five factors: the R&D budget size, the R&D budget distribution, the rate of technical learning, the technology development stage length, and production costs. The impact of the size of the R&D budget was examined by leaving the percent allocated per stage the same but varying the aggregate amount over the range of -50% to $+50\%$. The budget distribution was analyzed by reducing the budget for stages 4, 5, and 6. This particular scenario was designed to represent the status quo of the defense technology development process. Typically, funding for maturing a technology through the middle stages comes not from Science and Technology (S&T) organizations but through an acquisition effort in the technology development phase. To understand the influence of the rate of technical learning, the learning factor from the technical progress model was varied between 1 (no learning) and 2. In the basic experiment, all technology development stages are one year in length. To understand the impact of stage length, the scenarios were run with stage lengths of two years and three years. Finally, the influence of production costs was analyzed by varying the cost rate from -100% to $+100\%$ of the baseline value.

3.4 Results and Analysis

3.4.1 Results of the Basic Experiment

First, we will consider the results of the basic experiment. The average values of each of the output statistics are displayed in Table 3.1. Note that for compactness, system specific outputs are only shown for system 1. The results are similar for the other two systems. The most obvious question is how do these program outputs compare to real acquisition programs. As far as program duration, the distributional parameters for

concept development, system development, and production were derived from Fox, p. 29 [26] with an average program duration of 15 years. We see from Table 3.1 that the base case has an average duration of 14 years, which is fairly close. As for cost, Fox does not provide cost data, but a recent GAO report provides the cost and schedule performance of 62 current weapons system programs [37]. An analysis of these data reveals that the average program cost is approximately $16 billion. An important caveat is that these data cover a wide range of programs. Some are small upgrade programs that are short and inexpensive while others are major system of systems acquisitions that will take 30 years and cost hundreds of billions of dollars. Even so, we can see from Table 3.1 that the average program cost for the base case is approximately $16 billion. Thus, we can say that the simulation outputs are within the right order of magnitude for an "average" acquisition program.

Table 3.1: The average output values over 40 repetitions for the scenarios of the basic experiment.

Output	Base Case	Evolutionary	Revolutionary
Total Acquisition System Operating Cost ($ million, annualized)	5807	6410	5169
Capability Growth Rate (System 1)	0.16	0.179	0.138
Program Duration (System 1, years)	14.3	11.8	17.2
Program Cost (System 1, $ million)	16091	14668	16736

In order to understand these results fully, we will address each of the four outputs in turn. Figure 3.3a depicts the 95% confidence intervals for the average annual cost to operate the acquisition system. Clearly, evolutionary acquisition is the most expensive and revolutionary acquisition is the least expensive. If the technology policy is less aggressive with evolutionary acquisition, why would it be more expensive? To better understand this outcome, let us consider the average cost of the individual

programs.

Figure 3.3b shows the confidence intervals for the average program cost to acquire a system of type 1. Here we see that the average program cost is actually lower with evolutionary acquisition than revolutionary acquisition. So as evolutionary acquisition supporters suggest, using mature technology must lower program cost. Then why does the acquisition system cost more to operate under evolutionary acquisition?

The answer is revealed when we examine the average program duration or cycle time. In Figure 3.3c, we see that the program length is much shorter with evolutionary acquisition. With a shorter cycle time, acquisitions happen more frequently. Each cycle imposes overhead costs including system development, production, and deployment costs. Since these overhead costs are far greater than any savings that would result from more efficient management of the technology portfolio, the overall cost rises.

But does the additional cost of evolutionary acquisition buy the DoD anything? Figure 3.3d reveals that evolutionary acquisition results in a superior annual capability growth rate. The annual capability growth rate is the "average" annual rate of capability improvement. Much like an interest rate, even small differences in the rate can result in a huge difference in the level of deployed capability over the long-run. Thus, we see that there is a cost/performance trade-off governed by the technology maturity requirement. Allowing more immature technology hurts system performance because it takes longer to move technologies into the field, but since it incurs large production costs less often, it is also less expensive. Strictly enforcing maturity requirements, on the other hand, means shorter, less expensive programs that achieve high performance by moving technologies into the field more quickly. Unfortunately, this incurs production costs more frequently and results in increased operating costs for the acquisition system as a whole.
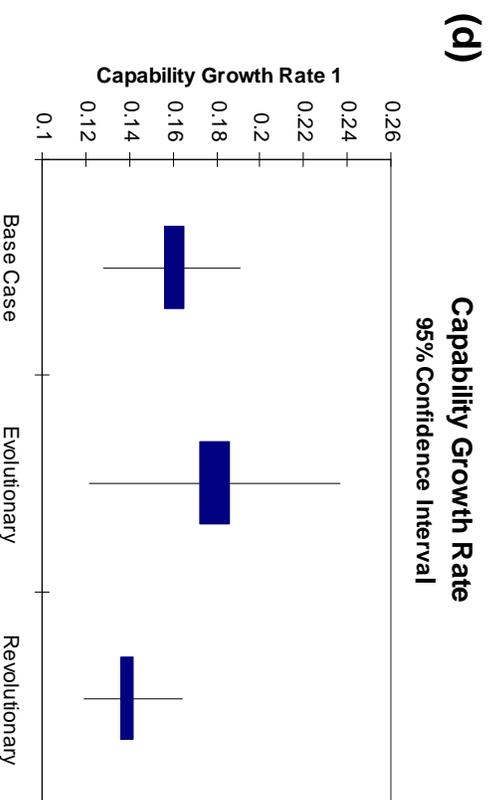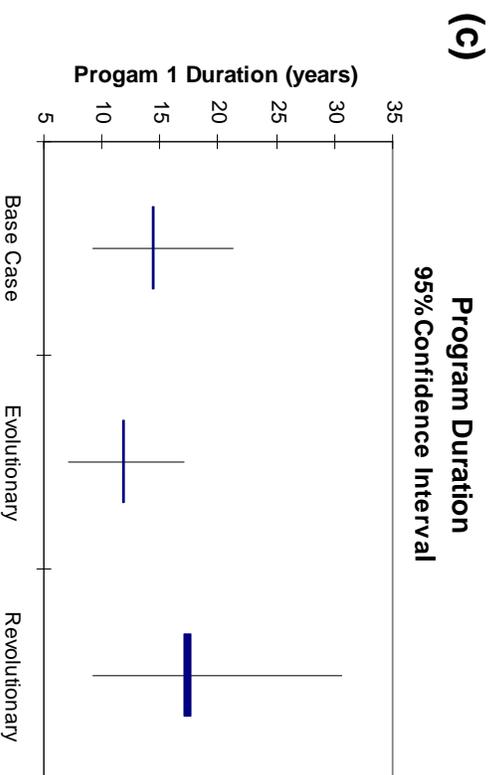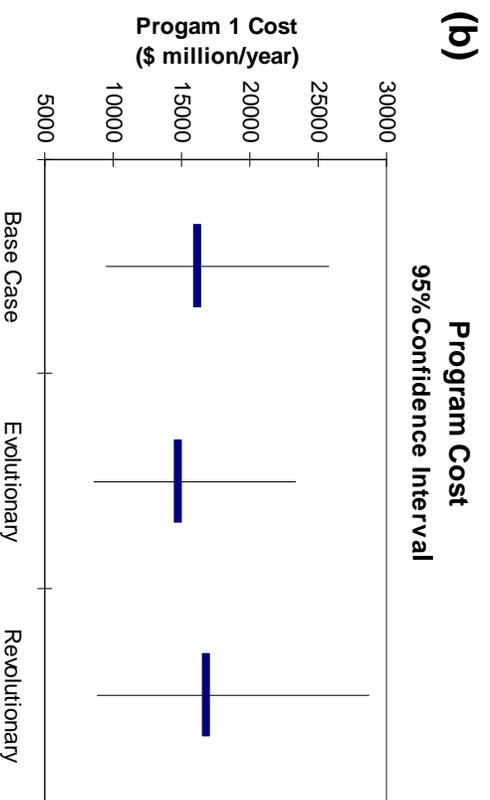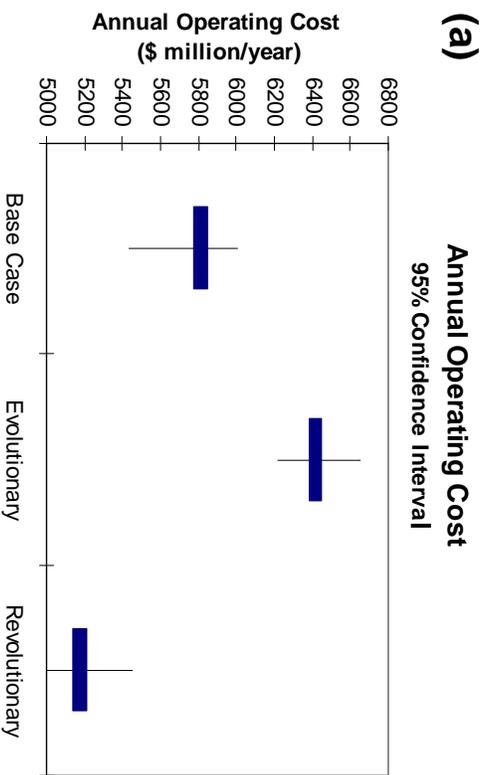
Figure 3.3: 95% confidence intervals by scenario for the mean values of each of the four primary outputs. The height of the box represents the width of the confidence interval, and the vertical lines represent the range of realized values.

This tradeoff is illustrated in Figure 3.4. In fact, by varying the technology policy one can move along a roughly linear frontier of cost/performance combinations. Figure 3.5 shows the cost and performance for all possible technology polices such that $1 \leq$ Min TRL ≤ 7 and Fallback TRL $\geq$ Min TRL. At first, this result would seem to suggest that technology policy should not be strictly enforced as budgetary restrictions would force changes in technology policy to meet cost goals. Fortunately, this is not the case.
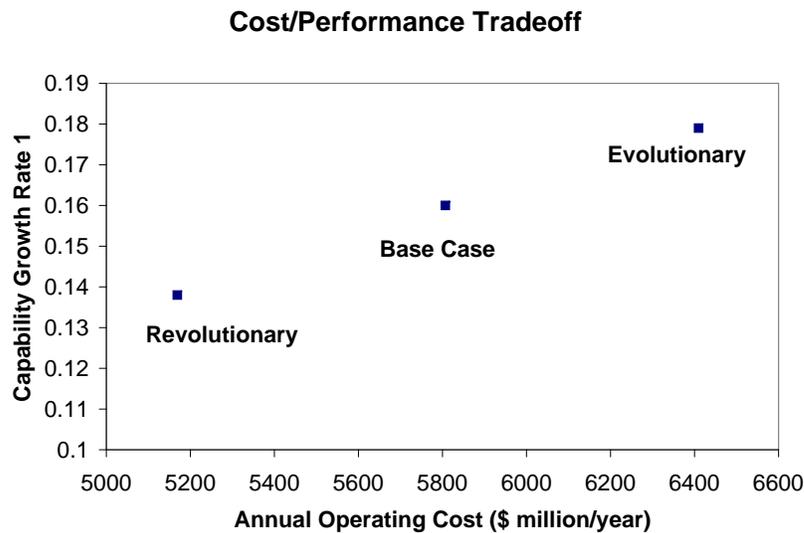
Figure 3.4: Cost/Performance trade-off for the basic experiment

In order to maintain a consistent, evolutionary technology policy but retain the ability to trade performance for cost, all that is required is to insert a delay between acquisition cycles. Figure 3.6 depicts the cost/performance combinations for the evolutionary policy with inter-cycle delays ranging from 0 to 7 years. Also shown is the linear trend line from Figure 3.5. Clearly, the introduction of a delay allows the evolutionary policy to replicate the cost/performance combinations achieved through shifts in technology policy. Thus, for any given cost target, an efficient policy can be found by imposing the evolutionary maturity requirements in combination with the appropriate inter-acquisition cycle delay.

Cost Performance Tradeoff
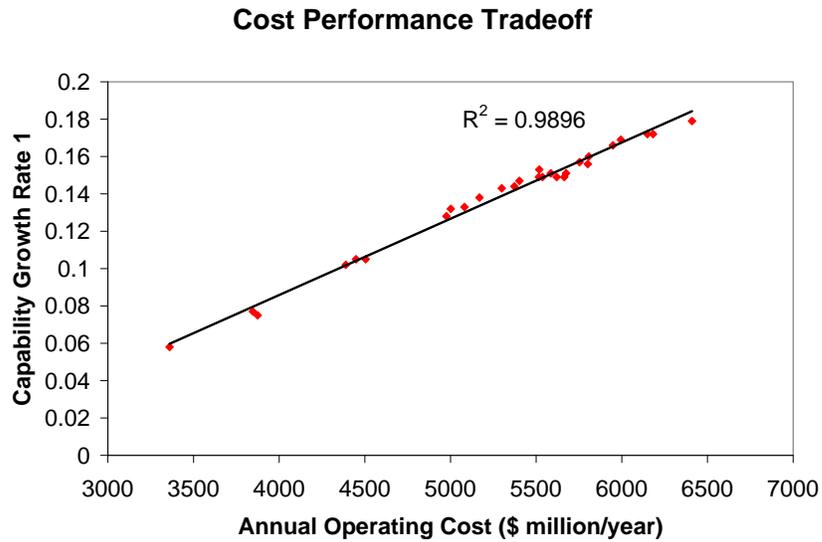
Figure 3.5: Cost/Performance trade-off for all possible technology policies with a linear trend line.

Cost Performance Tradeoff
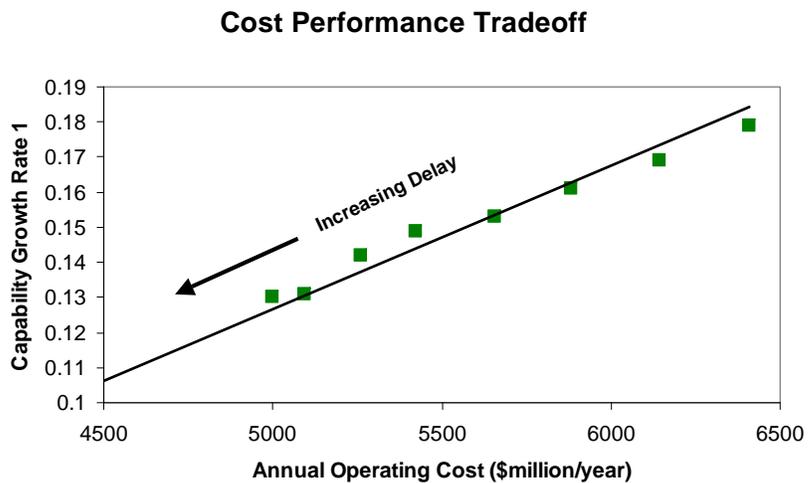
Figure 3.6: Cost/Performance trade-off replicated through the evolutionary policy with an inter-cycle delay

3.4.2 Sensitivity Results

The previous section presented the basic results of the experiment, but there remains a question of robustness. How stable are results? Are there any cases where the evolutionary policy is not the best performing? Through the five scenarios described in Section 3.3.3, we will examine the factors that ultimately drive the behavior of the system.

The first scenario we will consider relates to the size of the R&D budget. The size of the R&D budget impacts the rate that new technologies proceed through the development process. The R&D budget was varied over a range of -50% to $+50\%$ of the budget in the basic experiment. Obviously, this has an impact on the cost of operating the total acquisition system, but it was found that all policy scenarios were affected evenly. Where differentiation occurred was in the capability growth rate. In Figure 3.7, we see that for small budgets, all of the policies perform poorly. There are simply not enough new technologies moving through the the pipe to support a healthy growth in system capability. As the budget increases, however, the performance of the policies begins to diverge. When the R&D process is well funded, the evolutionary policy is clearly superior in terms of capability growth.

There is one other important point to note regarding the R&D budget. It is apparent from Figure 3.3d that the evolutionary policy has the greatest variability in the realized capability growth rate. This would seem to imply that while the evolutionary policy performs the best on average from a capability standpoint, it also appears to be the most risky. The higher variability is a result of the the evolutionary policy's dependence upon the R&D system. Unlike the other two policies, the evolutionary policy must make due with whatever technology is mature even if it is not a significant improvement on the existing capability. The other two policies, on the other hand, can fund a higher performing but immature technology to compensate. Thus, it would seem that the variability of the capability growth rate would be dependent

72

Capability Growth Rate vs R&D budget
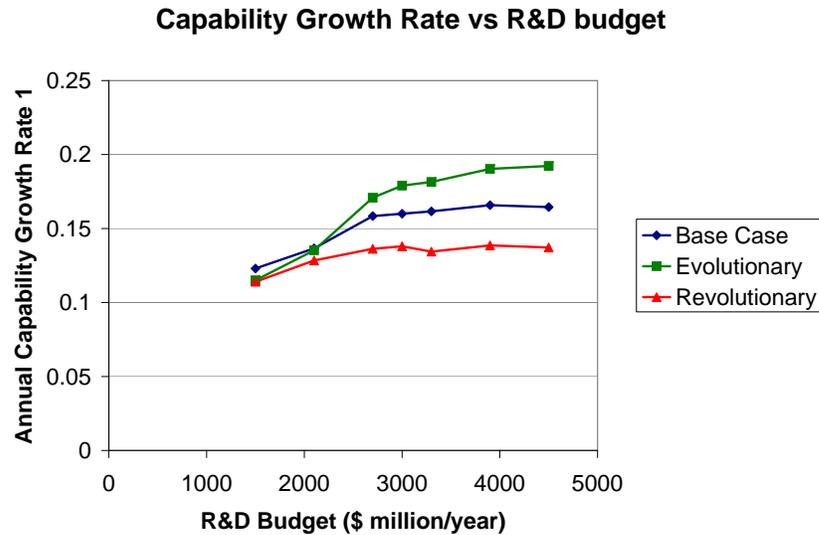
Figure 3.7: The annual capability growth rate versus the size of the R&D budget

upon the likelihood of high performing technologies being matured quickly. Figure 3.8 shows this to be the case. As the R&D budget increases, the standard deviation of the capability growth rate decreases rapidly for the evolutionary policy. The other two polices appear to be more robust. Thus, the riskiness of the evolutionary policy is highly dependent upon the health of the defense R&D system.

So what should the R&D budget be? Figure 3.9 displays the benefit/cost ratio (capability growth rate divided by the annual operating cost) versus the size of the budget. The results are quite revealing. First, we see that the policy that produces the most performance on the dollar is actually the revolutionary policy in combination with a 30% to 50% reduction in the R&D budget. Of course, this efficiency comes at the cost of fielded capability. The maximum benefit/cost ratio for the evolutionary policy is achieved by the original budget while the base case falls in between. Thus, pushing the budget beyond the $3 billion figure used in the basic experiment will increase performance and reduce risk but at a rising cost.

The next scenario also relates to the R&D budget, but in this case it is the distribution of the budget among the stages that is of interest. In particular, this

73

Std Dev of Capability Growth Rate vs R&D budget
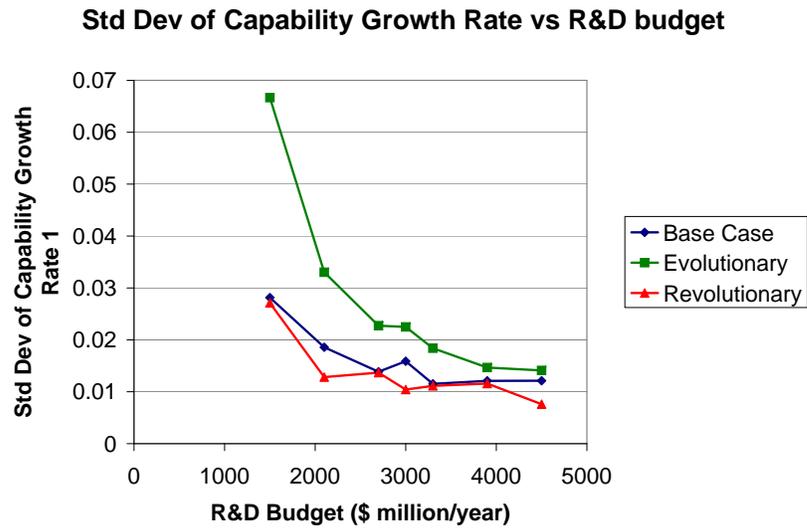
Figure 3.8: The standard deviation of the capability growth rate versus the size of the R&D budget

Benefit/Cost Ratio vs R&D budget
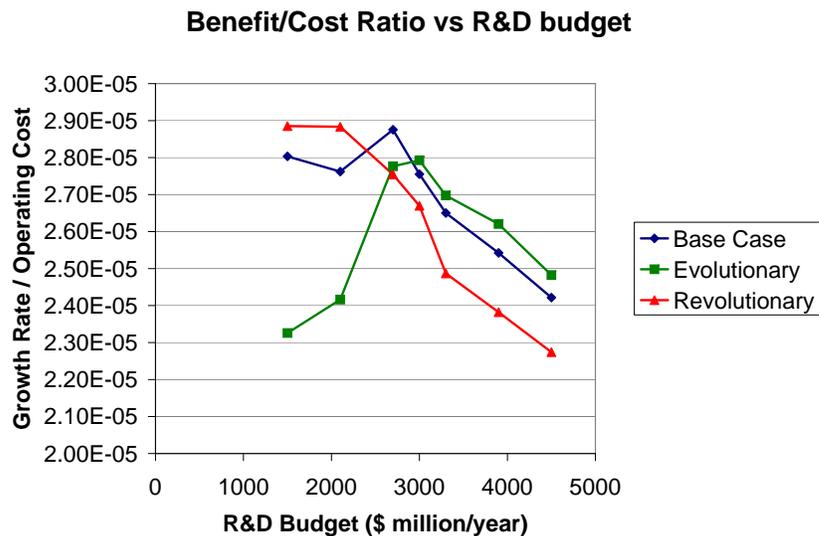
Figure 3.9: The ratio of the capability growth rate to annual operating cost versus the size of the R&D budget

scenario is designed to represent a situation that is often referred to as crossing the chasm. Crossing the chasm describes the difficulty that technology development efforts often encounter in moving through the middle stages of technology maturation because of a scarcity of funding. To simulate this scenario, funding for stages 4, 5, and 6 was varied over a range of 25% to 100% of the baseline value. Figure 3.10 reveals that the best policy from a performance standpoint is quite sensitive to the level of middle stage funding.
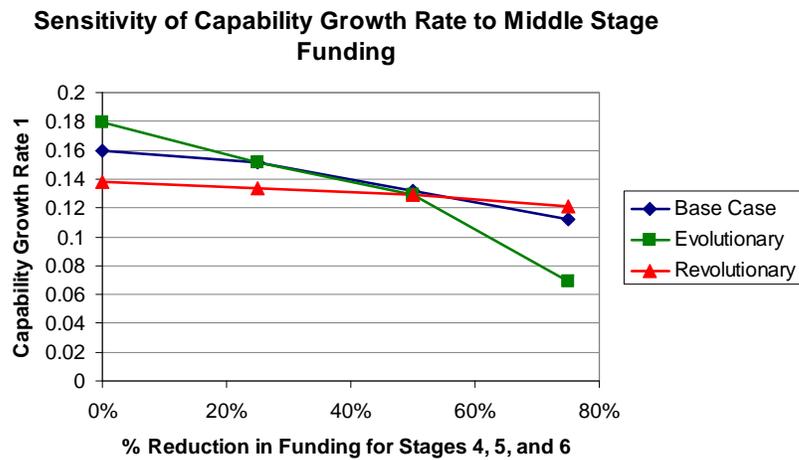
Figure 3.10: The capability growth rate when middle stage R&D funding is is cut

As we would expect, the evolutionary policy is the most sensitive since it is dependent upon a constant supply of mature technologies. On the opposite end, the revolutionary policy is the most robust since it can provide its own middle stage funding, and once again, the base case falls in between. Given the varied rates of performance decay among the three policies, there are domains where each is dominant. When R&D is well funded, the evolutionary policy provides superior performance. As middle stage funding is reduced by more than 25%, the performance of the base case policy begins to exceed the performance of the evolutionary policy. As funding declines further, the revolutionary policy becomes the top performing policy.

Of all of the scenarios presented in this chapter, the crossing the chasm scenario

is probably the most similar to business as usual at the DoD. Typically, S&T funding covers early stage technology development, but once technologies reach the middle stages, the only readily available source of funding is through an acquisition effort. The base case policy is also fairly similar to the risk mitigation strategy that many acquisition programs use: try to utilize the most promising technology, but if that fails, fall back to the existing, mature technology. Thus, it would seem that given the circumstances that most acquisition programs operate under, the business as usual policy is quite rational. Of course, it should be pointed out that all of the acquisition policies perform better when middle stage R&D is well funded.

Another critical factor that influences policy outcomes is the rate of technical learning. As was mentioned previously there are two components to the rate of technical progress in this model, the exogenous performance growth rate and the internal learning factor. When the exogenous growth rate is varied, all of the polices are affected equally because the exogenous technical progression occurs regardless of any actions taken in the defense acquisition system. When the learning factor is varied, however, the impact is significant.

Figure 3.11 shows the change in the capability growth rate for system 1 as the learning factor varies between 1 and 2. Clearly, the larger the learning factor, the greater the separation between the policies. It is apparent that the evolutionary policy achieves the superior capability growth rate as the learning factor increases. The driving force behind this behavior is the acquisition cycle time. Evolutionary acquisition exhibits the shortest acquisition cycle time, and thus, knowledge gained from fielding a system is accumulated more rapidly. The learning factor represents the impact that this knowledge has on new technology development. Consequently, as the value of learning increases, the gain from faster acquisition cycles increases.

The implication here is that evolutionary acquisition is more important for military specific technologies. When the military is the only user of technology, it is also

Capability Growth Rate vs Learning Factor
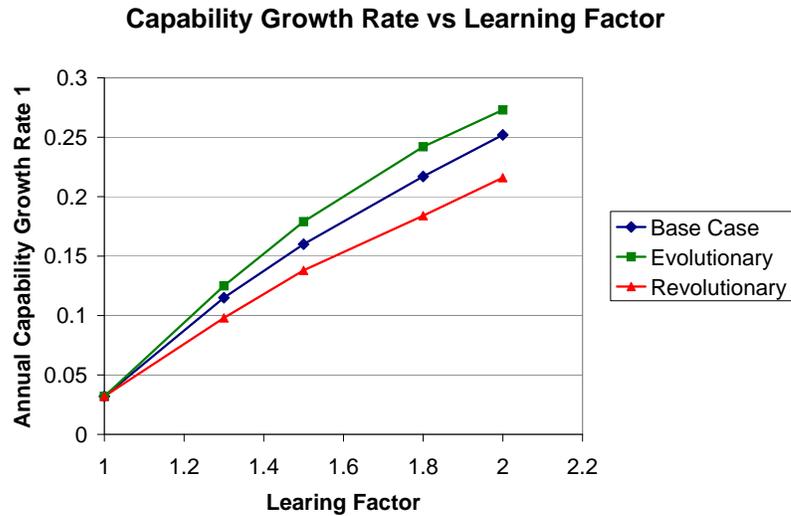
Figure 3.11: The capability growth rate versus the size of the learning factor

the only source of feedback to support future improvements. Consequently, the speed at which the feedback occurs affects the speed at which technology can improve. On the other end of the spectrum, Commercial Off The Shelf (COTS) technologies will likely improve regardless of the military actions. Thus, there is no real benefit to faster acquisition cycles from a performance improvement standpoint.

Another scenario that is of particular interest is the amount of time it takes to move a technology through each stage of the technology development process. In the basic experiment, stage length was deliberately deterministic and equal to one year for every stage. The rationale behind one year per stage was derived from a NASA study that determined that it takes about ten years on average to move through all nine TRL levels [45]. So one year per stage is a reasonable estimate for the average case. However, in reality, there is a great deal of variation in the maturation time of technologies. Some technologies mature extremely rapidly as in the semiconductor industry while others can take a very long time as in the pharmaceutical industry. It turns out that stage length has a major impact on the performance of the acquisition system.

In this scenario, the length of a stage in the technology development process was varied from 1 to 3 years. As Figure 3.12 demonstrates, the capability growth rate diminishes as the technology development cycle time increases. The key is the relationship between the acquisition cycle time and technology development time. When the time it takes to develop a technology increases beyond the length of a standard acquisition cycle, it means that the latest technologies are not making it through the process fast enough to be ready and mature for the next system under development. Thus, we are left with two policy choices. We can either select the most promising technology and finish maturation under the guise of an acquisition program, or we can simply utilize the same technology used in the last system that was deployed. In either case, the capability growth rate is diminished. In the first case, the acquisition cycle time is effectively extended as it must await the conclusion of a more lengthy technology development phase (Figure 3.13). This both delays deployment of the new capability as well as any learning that might occur from using the new technology in the field. In the second case, we maintain the same cycle time, but we in some sense deploy a new system that is identical in capability to the last. Effectively, the acquisition cycle time increases because it takes more cycles to get a new technology.

The final scenario represents the impact of production costs on the affordability of evolutionary acquisition. The production cost rate was varied from zero to $8 billion per year. Figure 3.14 reveals that as procurement cost increases the spread between the operating costs of the three policies increases. The shorter the acquisition cycle, the more frequently production costs are incurred and, consequently, the greater the impact of an increase in production costs. Conversely, the lower production costs are, the more cost effective evolutionary acquisition becomes.
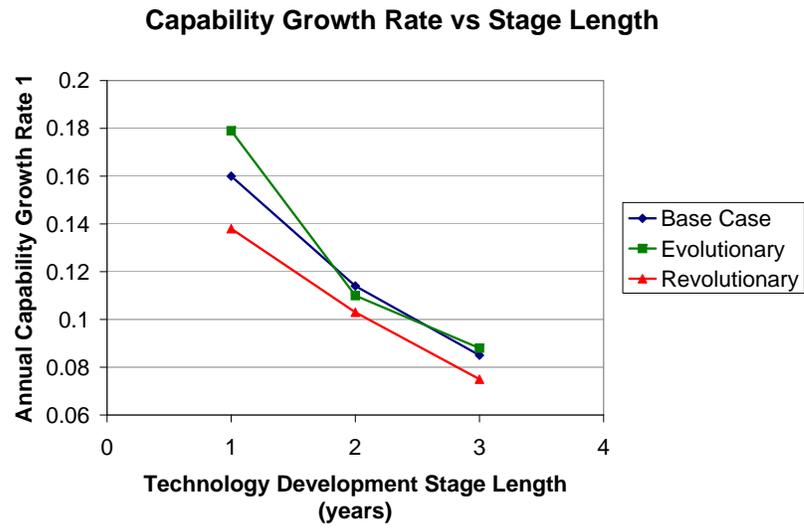
Figure 3.12: The average capability growth rate as a function of the time to complete a technology development stage
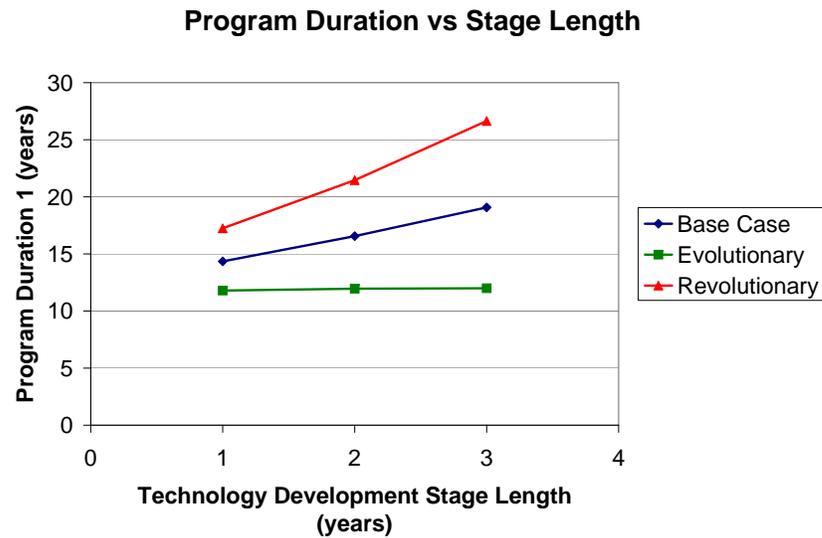
Figure 3.13: The average acquisition program duration as a function of the time to complete a technology development stage

Operating Cost vs Production Cost Rate
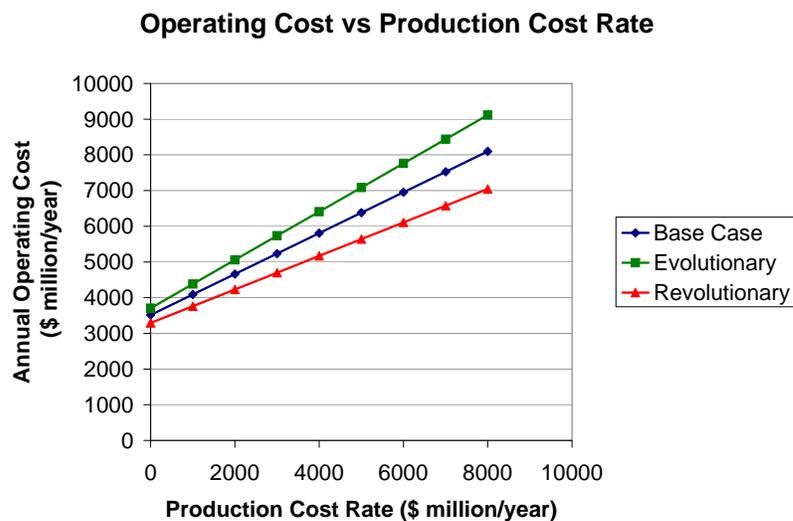
Figure 3.14: The annual acquisition system operating cost as a function of the production cost rate

3.5 Discussion

The production cost scenario raises several issues regarding evolutionary acquisition. Clearly, the more expensive it is to produce and deploy the next iteration of a system, the less affordable evolutionary acquisition becomes. But, of course, that is dependent upon the nature of the system under consideration, and this is a key difference between evolutionary practices in a commercial setting versus a defense setting. A commercial firm does not purchase its own product. In fact, if we take the example of a car manufacturer, there is always substantial portion of the customer base that is looking to buy a new car. Thus, the car manufacturer is going to build and sell cars continuously. The costs of upgrading a model might include the costs of any technology development, the cost of changing the design, and the cost of any retooling that must be done at production facilities. If the manufacturer is particularly successful, it may gain market share from its competitors, and thus, the investment pays for itself. Consequently, a commercial firm can actually make more money from cycling faster and using an evolutionary approach. When the DoD would like to buy a new weapon

system, it must pay for all of the same development costs plus it must purchase the product. Furthermore, if through more rapid acquisition cycles the DoD improves the performance growth rate of its systems, it may outperform its adversaries, but it does not generate a monetary return to help fund the faster pace of system development.

Thus, the cost of evolutionary acquisition is critically dependent upon the length and cost of stages in the system acquisition life-cycle. The simulation model presented in this chapter was generic in the sense in that it assumed that something was acquired in each cycle but it did not differentiate between say a new system design or a product upgrade. Representing either case could be achieved by simply changing the cost and duration parameters in the model. The key outcome of the evolutionary policy was that the acquisition cycle was shortened and the cost of each cycle was reduced simply by employing mature technology. In the examples above, however, the decline in cycle costs from more efficient technology development alone was not sufficient to compensate for the increase in the cycle rate. Thus, total acquisition costs rose with evolutionary acquisition. Some have suggested, however, that the length and cost of other phases of the acquisition life-cycle would decline under evolutionary acquisition as well. The idea is that if acquisition programs are less ambitious and shorter, development will be easier and there will be fewer problems with unstable funding. Thus, we should expect lower system development and procurement costs as well. Consequently, the question becomes, if the costs of system development and production decline under evolutionary acquisition, does evolutionary acquisition then become less expensive than more traditional methods?

To consider this question let us develop a very simple model for the cost of operating the defense acquisition system. First, we define the following symbols:

$r_{ij} \equiv$ the acquisition cycle rate for system i under policy j in cycles per year.

$C_{ij} \equiv$ the cost per acquisition cycle for system i under policy j.

$K_j \equiv$ the total cost per year for operating the defense R&D system under policy j.

$A_j \equiv$ the annual cost of operating the defense acquisition system under policy j.

We can define the cost rate to operate the acquisition system under policy j as

$$A_j = \sum_{i=1}^{n} r_{ij} C_{ij} + K_j$$

where n is the number of systems begin acquired. Thus, if policy e represents evolutionary acquisition and policy t represents traditional acquisition, then evolutionary acquisition would be less expensive if $A_e < A_t$. For the moment, let us assume that all systems being acquired have identical cost and cycle rates. This leaves us with the relationship

$$nr_e C_e + K_e < nr_t C_t + K_t.$$

Furthermore, if we assume that we keep our R&D budget fixed we can simplify even further to yield

$$\frac{C_e}{C_t} < \frac{r_t}{r_e}.$$

Of course, since the rate of acquisition is slower under the traditional acquisition policy, the right hand side will be strictly less then one. This implies that a simple decline in program costs from evolutionary acquisition is not sufficient to reduce the total cost to operate the acquisition system. Instead, program costs must to decline sufficiently to offset the increase in the rate of acquisition.

To better illustrate this point, imagine that the cycles were weekly and cost $10. The operating cost would be $10 per week. Now let us assume that we institute a new policy that reduces cycle costs to $8 per cycle but the cycles now occur twice as fast. That means that under the new policy the operating cost would be $16 per week. Thus, even though the cost per cycle decreased, the total cost increased.

When we consider defense acquisition cycles, if the system development and procurement costs also drop under evolutionary acquisition, that might seem to suggest

that we could overcome this deficit. If, however, the durations of system development and technology development also decrease, then the equivalent cost threshold becomes even more difficult to reach. Furthermore, if we consider spiral development where there are several short, overlapping cycles, we see that we would require fairly low development, production, and deployment costs to compensate for the speed of the cycles.

Thus, the critical question becomes, how does evolutionary acquisition affect the length and cost of development and procurement activities versus a traditional single-step to capability approach? This is not a trivial question, and the answer will likely depend on the type of system being acquired. Upgrades to complex, integrated systems can lead to substantial design modifications to accommodate even seemingly simple changes and using more mature technologies does not correlate to easier integration [71]. In fact, experiences at Westland Helicopters indicate that even when a system such as a military helicopter is designed with modularity and upgradability in mind, changes can unexpectedly propagate through large portions of the system design [12, 24]. At the other end of the spectrum, systems with very loose coupling between system components may be quite amenable to rapid upgrade and change. Perhaps the most extreme example of this type of system is the Internet where the system architecture changes continuously without any supervision or control.

Thus, this issue merits substantial additional research and is really the determining factor regarding evolutionary acquisition's potential for cost savings. This is not to suggest that if the costs of acquiring a particular system type do not decline under evolutionary acquisition that the approach is useless. The results of this study suggest that evolutionary acquisition delivers other benefits such as a boost in the capability of systems actually deployed in the field. Instead, it simply means that additional capability will continue to come at additional cost. Consequently, cost and performance may be traded off by simply appropriately spacing acquisition cycles.

3.6 Conclusions and Further Research

The results from this simulation study lead to some highly suggestive findings and critical avenues for future research. First and foremost, even with a first-order representation of the acquisition system, the results suggest that the adoption of evolutionary acquisition policies has the potential to improve the performance of deployed systems. However, lower operating costs for the defense acquisition system are not automatic. While each individual program should be less expensive under evolutionary acquisition policies, the faster acquisition cycle time means that development, production, and deployment costs are incurred more frequently. This may overwhelm any cost savings from managing technology development more efficiently. As discussed in Section 3.5, these cycle costs must decline sufficiently under evolutionary acquisition to achieve net cost savings. Thus, depending on the type of system being acquired, evolutionary acquisition may actually be more expensive than traditional means of acquiring military systems. This is a critical issue for future research. However, this should not be interpreted as an endorsement of traditional acquisition methods. Instead, acquisition cycle time can be used to control the costs of an evolutionary policy without reverting to a traditional approach that employs immature technology. A requirement for mature technologies can be consistently imposed with the next acquisition cycle beginning only when it is affordable.

There are some important caveats on this conclusion, however. First, the above results are more significant for military specific technologies than commercial technologies. Commercial technologies will continue to develop and improve regardless of the actions of the DoD because the DoD is actually a small player in the market. One example is microprocessor technology. On the Commanche helicopter program, the mission processing technology was changed three times because Intel introduced newer processor models faster than the DoD could develop an advanced combat helicopter [61]. For military specific technologies, however, forward progress is dependent

84

upon actually testing and fielding a technology and gathering user feedback. Thus, the faster acquisition cycles are, the faster learning can be incorporated into new technologies under development. Of course, faster acquisition cycles also mean that exogenously developed commercial technologies can also be moved into the field faster.

Second, evolutionary acquisition policies do not function well when the R&D process is underfunded. Evolutionary acquisition depends on a steady stream of mature technologies. When the research pipeline is "starved", not only does the performance of deployed systems decline on average, but it also becomes more unpredictable. More traditional acquisition methods mitigate this this risk by using an acquisition effort to secure funding for technology development.

Third, the underfunding of middle stage technologies, as is typical for government technology development [14], also adversely impacts evolutionary acquisition policies. Under these circumstances, traditional approaches to acquisition are actually superior to evolutionary methods since they mitigate the risk of technologies failing to cross the chasm. Thus, it would seem that business as usual is quite reasonable under the current funding environment for military R&D activities. Though, it is important to point out that traditional acquisition policies under this scenario still underperform evolutionary policies when R&D is fully funded.

Fourth, the relationship between the time required to develop a technology and the acquisition cycle time is crucial. Essentially, the pace of technology development dictates the pace of acquisition. When technology development is slow, acquisition must slow down to accommodate. Thus, if the acquisition cycle time is already close to the technology development cycle time, there may be little, if any, advantage to shortening the acquisition cycle time through the application of evolutionary acquisition policies.

Finally, there are several features of the current defense acquisition system that were not considered in this analysis. First and foremost among these is concurrency.

For major acquisition efforts there is often substantial overlap between the technology development, system development, and production phases. While this is often an attempt to compress an otherwise long acquisition cycle, the resulting rework often increases costs and leads to performance shortfalls. This problem has been extensively documented elsewhere, and there is no need for it to be recapitulated here. If, however, the imposition of evolutionary acquisition and its shorter acquisition cycles reduced the temptation to use a concurrent acquisition strategy, it is possible that there could be a net cost savings through the reduction of rework, but that determination must be relegated to future work. Other features of defense acquisition not considered in this model are operations and maintenance costs, basic research funding, non-centralized acquisition management, program cancellation, program budgeting, the capacity of the industrial base, the capacity of the government to consume, and system integration issues. Each of these factors certainly influence the behavior and cost effectiveness of the defense acquisition system and may be examined in future work.

What we can ultimately derive from this study is that, at least to a first order, there are definite benefits to the better management and development of new technologies implied in evolutionary acquisition. The outstanding question raised is whether or not there is a net reduction in cost when we consider the entire acquisition system, not just a single program. What this study revealed is that net cost savings are not automatic, and additional research is required to determine under what circumstances they are possible. Furthermore, when we consider acquisition policy in general, this study reveals the importance of considering the entire system when evaluating a policy. As in Chapter 2, we see that the implications of a reform were not fully understood by those who advocated it. Only by treating the defense acquisition enterprise as a system can one hope to understand the implications of alternative acquisition policies. The model presented in this chapter provides one

example of how this type of problem can be approached in a systematic fashion to inform policy decisions.

CHAPTER IV

A METHOD FOR VALUING DEFENSE ACQUISITION PROCESS IMPROVEMENTS

While Chapters 2 and 3 primarily dealt with evaluating the most recent acquisition reform initiative, evolutionary acquisition, this chapter deals with the more general problem of quantitatively valuing changes to the defense acquisition system. As indicated in Chapter 1, there is often no means to objectively value changes to the defense acquisition system. This is because national defense is a public good, and there is no market price to treat as a consensus value. Thus, it becomes quite challenging to objectively compare and trade off among acquisition reform policy options.

When monetary valuation is not possible, utility theory is often a convenient means of evaluating policy alternatives. Unfortunately, utility theory is difficult to apply in the defense acquisition context since there is such a diverse set of stakeholders. Consequently, the traditional approach to valuing improvements to an acquisition process is to assume that the quantity and type of systems acquired remain fixed. Thus, there is no change in military utility. In that case, only costs change, and a Net Present Value (NPV) analysis can be performed on the resulting cost savings.

This approach has three major shortcomings. First, cost savings will lead to changes in how many systems are purchased and/or what type are purchased. Thus, the change in value is more than just cost savings. Second, risky initiatives may be staged, and NPV analysis fails to account for resulting downside loss mitigation. Consequently, NPV may significantly understate the value of reform opportunities. Third, there is no systematic means to account for risk aversion.

In this chapter, a method is presented to address all three of these issues[1]. First, a pricing index will be used to measure changes in buying power. Conceptually analogous to measuring inflation, the use of a pricing index allows us to capture monetarily changes in what is acquired. This monetary valuation allows modern financial analysis tools to be applied to what is effectively an investment opportunity. Options analysis, in particular, will be used to properly account for the risk mitigation inherent in a staged investment. Finally, to allow decision-makers to trade-off risk and return, a risk/return portfolio method is developed to assist risk-averse decision-makers.

To illustrate these methods, they are applied to a notional process change in the area of military shipbuilding. What was found was that a traditional approach to valuing acquisition process improvements can, in some circumstances, significantly understate value and lead to rejection of an otherwise valuable opportunity.

4.1 The Valuation of Acquisition Reform

How does one value national defense? What is the value of adding one more aircraft carrier to the fleet? The answers to questions such as these probably depend upon the person asked, and there is likely a great deal of variation among individuals. Since changes to the defense acquisition system are ultimately intended to improve the quality of national defense provided by the US military, an inability to value national defense would seem to inhibit a systematic means to evaluate defense policy alternatives. And, as Cancian noted, there is decided lack of objective means to evaluate defense outcomes [9]. All is not lost, however, because transforming the defense acquisition enterprise is really about efficiency. The question is not what should be bought, but instead, how can the government buy systems with highest possible quality at the lowest possible cost.

---

[1]An earlier version of the ship production model contained in this chapter was presented in Pennock, et. al., 2007 [57]. That model was a special case of the more general model presented here.

If we assume that the US defense policy apparatus is at least reasonably capable at determining which military systems are needed (and there is some evidence to suggest that it is [41]), then the evaluation of acquisition reform policy is somewhat simplified. Since demand always exceeds the available resources, it is not necessary to assess the ultimate value of national defense realized by a particular policy. Instead, one only needs to determine which policy allows the acquisition system to meet more needs within the available resource constraints. This is in contrast to the somewhat dubious method of tying military value to market comparables advocated by Housel and colleagues [42].

At first, it might seem that simple cost savings would be sufficient. Unfortunately, simple cost savings as a measure of value implicitly assumes that the government is a price taker in the market for military systems. This is certainly not the case. The market for military systems, at least in the United States, is effectively a monopsony. Consequently, virtually every decision that the US government makes regarding the acquisition of military systems affects the health and well-being of the defense industry. In particular, the US defense industry maintains a significant excess of production capacity. This is intentionally encouraged by the government to provide surge capability in the event of a major war [28,66]. This is compounded by the fact that most military systems are acquired in short, intermittent production runs. The combination of these two factors means that defense production exhibits increasing returns to scale.

The most immediate consequence of increasing returns to scale is the compounding effect of cost overruns. Nearly every acquisition program experiences cost overruns due to the systematic underestimation of costs by both industry and government [2,10, 23]. Thus, over the course of acquisition programs, the per unit cost of systems such as aircraft and warships typically rise far higher than what was originally planned. As a result, the available budget is often insufficient to pay for the entire planned

production run, and the size of the production run is cut. Unfortunately, in an increasing returns to scale situation, cutting the number produced actually increases the per unit cost even further. This means that the production run must be reduced even further. Thus, an acquisition program is hurt not only by the initial cost overrun itself but also by the subsequent adverse impacts on production efficiency. This is exactly the situation that programs such as the F-22 have experienced [35].

The flip side to increasing returns to scale is that cost savings trigger a compounding effect in benefits. Thus, when program costs decline and free up more resources, the size of the production run can be increased or at least remain closer to the intended size. This means that the overhead associated with the excess productive capacity can be shared over a greater number of units, and it allows producers the opportunity to invest in new plant equipment and more efficient means of production. Thus, the per unit production costs actually drop, and the compounding effect means that the size of the production run can be increased even further. This is the phenomenon that simple cost savings fails to capture. Thus, the net effect of a change to the defense acquisition system is a change in buying power, and buying power is measured via price indices.

4.2   Price Indices

A price index is a way of assessing the value of a bundle of goods and services. As the price of goods and services change over time, the cost of the bundle changes. An increase in the cost of the bundle indicates a decrease in buying power since a consumer would be able to afford fewer instances of the bundle. In the converse, a drop in cost means an increase in buying power. While price indices are normally used to measure inflation, they have a logical analog in the valuation of defense acquisition reforms.

To illustrate this point, let us consider the case of deflation. Deflation occurs when

there is a general decline in the price level of goods and services. Let us assume that there there is a decline in a price index between two time periods. That means that under a fixed budget one could purchase more goods and services in the second time period than the first. Thus, buying power has increased. One should note however, that the same buying power could also be achieved by increasing the budget in the first period such that an equivalent amount of goods and services can be purchased in the first and second periods. Thus, we could say that the value of the deflation is the difference between this increased budget and the original budget. One would be indifferent between having either the increased budget or deflation. In the abstract, there is no difference between deflation and an increase in the government's buying power through acquisition reform. The only difference is that deflation measures a change in buying power over two time periods while the government's buying power is measured over two scenarios, acquisition with the reform and acquisition without the reform. This concept will be explored further in Section 4.5.

There is a well-developed body of theory in economics regarding the use of price indices. For the purposes of this exposition, however, only its most basic tenants are required. Probably two of the most common price indices are the Laspeyres and Paasche indices [58].

$$
\begin{aligned}
P_L &= \frac{p_1 \cdot x_0}{p_0 \cdot x_0}, \\
P_P &= \frac{p_1 \cdot x_1}{p_0 \cdot x_1}.
\end{aligned}
$$

These indices assess the change in the cost of a pre-specified bundle of goods, but in reverse time order. The Laspeyres index assesses the cost of buying a bundle of goods from the starting time in a future time while the Paasche index considers the cost of buying the bundle of goods from the end time in a previous time. The two bundles may differ because of substitution and income effects. This means that the two price indices can differ. The Fisher price index [58] attempts to split the difference by

taking the geometric mean of the two

$$P_F = \sqrt{P_L P_P}.$$

The substitution issue merits some additional discussion. When consumers face relative price shifts in goods and services, they may alter the composition of the bundle that they purchase. For example, if the price of natural gas increases and the price of electricity declines, consumers may substitute electrical heating for natural gas heating. Likewise, in the defense context, a shift in the cost of one military system may prompt a shift in the force structure. For example, if the cost of air superiority fighters rises relative to the cost of surface-to-air missiles, the military may opt to substitute surface-to-air missiles for fighters, since there is some overlap in their role. Practically speaking however, this is unlikely in most circumstances. Given the nature and momentum of acquisition programs, it is unlikely that any cost savings realized in one program will be transferred to another, at least in the short-term and for relatively small improvements. For that reason, the analysis presented in this chapter will assume that there is no substitution effect or income effect, and consequently, all three of these price indices will be identical. Be that as it may, it would be difficult to deny that substitution can and does occur over the long-run for large changes in the acquisition enterprise. When substitution is a concern, a more sophisticated index such as the Konus index may be used, but doing so would require an extensive assessment of how policy makers would alter the force structure in the event of major cost shifts.

There is one final caveat to note regarding the use of price indices in this analysis. The Laspeyres and Paasche price indices do not account for changes in the quality of goods procured. For example, the cost of a good such as an automobile may increase, but the quality may also increase. Thus, consumers may not decrease their consumption of automobiles because they are getting more for their money. Since it is generally agreed that there have been dramatic increases in the quality of US

military systems since World War II, it might at first seem that this would be an issue. Fortunately, with the assumption of sticky substitution and the tendency of the military to maximize performance regardless of cost, it is not. If an acquisition process improvement allows the DoD to increase procurement quantities, it simply means that the DoD will be able to procure more systems at each increase in quality level than it would have previously. If quality based substitution is an issue however, it would have to be addressed using a more sophisticated price index such as the Konus index mentioned earlier.

4.3 Applying Investment Analysis to Defense Acquisition

If the application of a price index allows one to assess shifts in buying power monetarily, then, logically, changes in buying power over time are effectively a cash flow stream. Of course, cash flow streams form the basis of investment analysis. Thus, the question then becomes how much should one be willing to pay in order to obtain a particular cash flow stream? If the value of a cash flow stream exceeds the cost to obtain it, then it is a worthwhile investment. Applied to the defense context, the value of the increase in buying power from an acquisition process improvement must exceed the cost of implementing it.

4.3.1 Net Present Value Analysis

As mentioned previously, the traditional approach to valuing cost savings in the DoD is Net Present Value (NPV) analysis. NPV is calculated by discounting benefits and subtracting the discounted costs. Under traditional capital budgeting, when NPV is applied to a commercial investment opportunity, the discount rate should be the cost of capital of the firm [54]. More commonly, however, the discount rate is set to an arbitrary hurdle rate felt to be commensurate with the level of risk. Risk is really the driving force in the valuation of investments. If risk were not an issue, investment valuation would be trivial, as one would only have to account for the time value of

money. In reality, however, there can be a great deal of uncertainty in realized cash flows.

In the context of this analysis there are really two types of risk: market risk and technical risk. Market risk is the uncertainty that results from random fluctuations in the marketplace. If one were considering investing in an exchange traded stock, market risk would be the uncertainty in the future price of the stock. In the context of defense acquisition, market risk would be the uncertainties in the prices of inputs that are purchased in the course of developing and producing military systems. For example, uncertainties in the price of steel or in the labor costs of shipyard workers would constitute market risks for defense acquisition.

Technical risk is the uncertainty in the execution of a project. For example, in a commercial setting, technical risk might manifest itself as the uncertainty regarding the efficiency of a new manufacturing process. It is quite similar in the defense acquisition context. Budgets and priorities at the DoD can shift from year to year, and changes in appropriations that may seem insignificant to Congress may seriously imperil a program [26]. Beyond programmatic issues, there is always a risk that the process improvement idea simply does not work. Intuitively, as the level of market or technical risk increases, an investment becomes less appealing. Greater risk requires greater return. In an NPV analysis, this is captured through the discount rate.

For simple now or never decisions, NPV is a perfectly adequate tool. The problem arises when NPV is applied to investments that occur over time and present opportunities to change course. When an investment contains embedded options, NPV may undervalue that investment. Since most real, as opposed to financial, investments occur over time, decision-makers often have the opportunity to terminate, ramp up, scale down, or otherwise alter the course of the investment based on new information that becomes available over time. For example, if during the development of a new product there is a significant technical failure, it is likely that management will

terminate the project. While the money already invested is lost, future losses from a failed product are avoided. Thus, the option to terminate effectively limits downside risk. In fact, most complex projects are staged for this reason. The shortcoming of NPV as an investment evaluation tool is that it does not account for such options. It assumes that once an investment decision is made, there is no turning back no matter how dismal the failure. For that reason, NPV has been heavily criticized as a business decision making tool, and it has been accused of leading to significant underinvestment [39, 40].

4.3.2   Options Analysis

In many cases, innovative approaches to acquisition are implemented in stages. Take, for instance, the case of the Arsenal Ship. The program to develop the Arsenal Ship employed several unorthodox approaches to ship acquisition including setting only a few broad performance goals, giving design responsibility to the contractor teams, and setting affordability as a requirement [46]. While the Arsenal Ship was ultimately canceled, it provided an opportunity to test new approaches to warship acquisition without jeopardizing the entire US Navy shipbuilding program. Thus, it would be naïve to employ NPV to value acquisition process improvements when the DoD attempts to mitigate the risk of new initiatives through testing and staging.

To properly account for such staging, options analysis is required. Sometimes referred to as real options, the options approach employs the stock option as its motivating metaphor. Consider a call option. Purchasing a call option gives the option holder the right, but not the obligation, to purchase shares of a stock at predetermined price. The holder would never exercise the option at a loss, and this feature of the contract limits exposure to downside risk. In a similar fashion, applying innovative acquisition methods to the Arsenal Ship program gave the government the right, but not the obligation, to employ those same methods on future warship

acquisition programs.

Options analysis traces its origins to the seminal paper by Black and Scholes that presented a closed-form equation for the price of a European call option [5]. Since then, an entire profession has evolved around the pricing of options contracts. In all cases, options derive their value from the behavior of an underlying asset, hence the often used term, derivative. The term "real options" was coined by Stewart Myers [53] in recognition of the similarities between many real investment opportunities and financial options contracts. In the case of real options, the underlying asset is not traded in financial markets. Examples include a natural resource such as an oil reserve, a production asset such as a factory, or intellectual property such as a patent. The canonical example of a real option is a lease on a petroleum reserve. The lease provides the holder the option to drill and extract oil if it proves profitable to do so. Of course, the value of real assets are subject to market fluctuations just like any other economic variable, and if the stochastic fluctuations in value can be replicated through a portfolio of market traded assets, then an option on a real asset can be valued just like a financial option.

Finding the value of an option is tantamount to solving a dynamic programming problem. The key issue is the discount rate. When an investment problem is solved using a traditional dynamic program, the discount rate is specified exogenously. As indicated in Section 4.3.1, the appropriate discount rate is dependent upon the level of risk, but risk changes with time and actions taken. Options analysis avoids this problem through the use of a replicating portfolio of market traded assets. The value of the replicating portfolio implicitly determines the discount rate. According to modern investment theory, this implied discount rate is consistent with the risk aversion exhibited by the shareholders of a publicly traded firm.

More specifically, shareholders require compensation for systematic risk. Systematic risk is risk that is inherent to the marketplace and cannot be eliminated through

diversification of investments. The higher the level of systematic risk exhibited by an investment, the greater the rate of return required by shareholders or, equivalently, the greater the discount rate. Technical risk would be non-systematic, and thus, shareholders would not require an adjustment to the discount rate to compensate for technical risk. They could simply diversify this risk away by holding other investments in their portfolio. Market risk, on the other hand, is part systematic and part non-systematic. The level of systematic risk is determined by the degree to which price movements of the underlying market-traded asset correlate with the market as a whole. The more correlated an asset is to the entire market, the more non-diversifiable risk it introduces into an investor's portfolio. As a result, risk-averse investors demand a greater return to compensate. If, on the other hand, an investor is risk-neutral, he or she is indifferent to risk, and the risk-free rate of return is a sufficient discount rate.

Considering the nature of the defense acquisition enterprise, it would be questionable to extend the shareholder metaphor to government decision-makers. There is no extant market with which to compare or evaluate defense acquisition reform efforts. Furthermore, with the vast resources of the Federal government, acquisition decision-makers should theoretically be risk-neutral, though this is not likely true in practice. Thus, for the purposes of this analysis, defense acquisition process improvements are evaluated as options but assuming risk-neutral decision-makers. This means that valuation only requires the risk-free discount rate. While this in some sense reduces the problem to a traditional dynamic program, the options metaphor is convenient, and the valuation method developed in this chapter utilizes the mathematics and tools developed to evaluate options. The risk attitudes of acquisition decision-makers are handled separately through the portfolio approach described in the Section 4.7. For a more extensive exposition on options analysis applied to real investment opportunities see Dixit and Pindyck [17].

4.4 A Model of Military Ship Acquisition

The concepts of applying price indices and options analysis to value acquisition process improvements have thus far only been presented in the abstract. In order to make them more concrete, they will be applied in the context of military ship acquisition. Since the start of the Cold War, the growing cost of US Navy warships has outpaced inflation and posed serious challenges to meeting the Navy's force structure goals [1]. Since the cost of an individual ship grows faster than the shipbuilding budget, each year the Navy can afford to procure fewer and fewer ships. While the decline in numbers has been partially offset by the rapid increase in the quality or capability of Navy ships, the situation is not sustainable, and if nothing changes, eventually the Navy will not be able to afford any ships at all.

Many have suggested that military shipyards are inefficient, and the application of commercial shipbuilding techniques could substantially reduce costs. According to a RAND study, however, labor, material, and equipment costs only account for about half of the cost growth in ships and have been roughly in line with inflation [1]. The remaining cost growth is attributed to customer-driven factors. Furthermore, another study points out several key differences between naval and commercial ships including [4]:

- Commercial ships are large and simple while military ships are relatively small and complex.

- The process for buying a commercial ship is much simpler than the government's approach to buying a warship.

- Design and construction of commercial ships is much simpler. They are essentially large steel boxes while warships are very complex with a high density of integrated, sophisticated equipment.

- Military shipbuilding employs much more engineering support than commercial shipyards. This results in a more expensive workforce.

Because of these differences, it is likely that commercial shipbuilding practices will not yield the same levels of efficiency that they do in commercial yards. For example, the hull cost is a much smaller percentage of the total cost for a warship than a commercial ship. Most of the cost is driven by the equipment installed on the ship rather than the ship itself. This would suggest that any efficiency improvements in ship construction will result in a smaller percentage cost savings than would be realized in a commercial context. Instead, additional cost savings must come from elsewhere in the ship acquisition enterprise. As Figure 4.1 indicates, there is more to the enterprise of shipbuilding than just the shipyards.
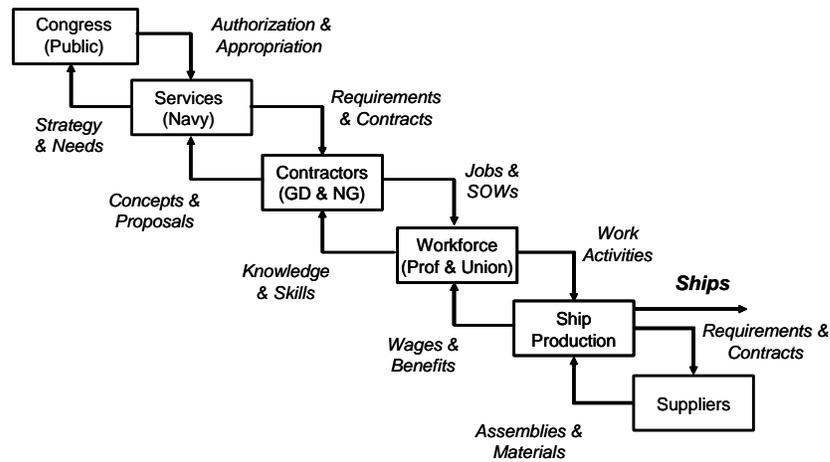
Figure 4.1: The US Navy Shipbuilding Enterprise [57].

As indicated by Rouse, enterprise transformation is driven by experienced or expected value deficiencies, and is enabled by changes in work processes [62–64]. Clearly, the Navy is experiencing a value deficiency, but what are the work processes that should be changed? A consideration of the ship acquisition enterprise suggests two categories that are candidates for reform: organizational processes and technical processes. Organizational processes consist of the methods by which the government

monitors, controls, and executes acquisition. These would include processes for authorization, appropriation, development, procurement, and deployment. Technical processes are related to the design, production, operation, and maintenance of ships. For our notional example, we will assume that the Navy has proposed several changes that will streamline the development and design process and reduce rework. In order to evaluate the impact of these process changes, we will require a model of ship production.

To that end, let us define the following:

$B(t) \equiv$ rate of cash flow from the shipbuilding budget at time t,

$X(t) \equiv$ rate of consumption of shipbuilding inputs at time t,

$C(t) \equiv$ cost of shipbuilding inputs at time t,

$Y(t) \equiv$ rate of ship production at time t.

Several assumptions were made to maintain simplicity and interpretably. These are summarized in Table 4.1. First, we will assume that all state variables are continuous. Since we are concerned with the long-run effects of process changes, short-run discontinuities will have a minimal impact. As for building the ships themselves, it is assumed that ship construction requires only one type of input. This is not an inherent limitation of the approach, but rather, this is merely to avoid complicating this notional example. Substitutability of inputs would require consideration in a more detailed analysis, but here we will assume that required inputs such as labor and materials are used in fixed proportions. Thus, we can treat all inputs as a single package.

Next, we must model the cost of our input package. As noted earlier ship costs have risen exponentially, and this is due to several factors including shipyard costs and increasing complexity. Since these trends show no sign of abating, it is fairly safe

Table 4.1: Summary of Modeling Assumptions.

Assumption	Justification
Shipbuilding budget follows geometric Brownian motion.	Allows us to consider the impact of growing or declining volatile budgets.
Ship construction input costs follow geometric Brownian motion.	Models the exponential growth in ship cost while accounting for economic noise in prices.
Ship production process is governed by a Cobb-Douglas production function.	Allows us to consider the impact of economies of scale on the quantity of ships produced.
Ship production is continuous.	Allows us to focus on the long-term trends in ship production sustainability.

to assume that input costs for ships will continue to rise exponentially. Of course, input costs are governed by economic forces, and consequently we would expect the cost of our input package to fluctuate in price over time. Geometric Brownian motion is a standard way to model prices that grow exponentially, so the cost of the input package will be modeled with the following stochastic differential equation,

$$dC = \alpha_C C dt + \sigma_C C dZ,$$

where α_C is the expected growth rate of input cost, σ_C is the volatility of input cost, and dZ is an increment of standard Brownian motion (i.e., a Wiener process).

Next, we must consider the shipbuilding budget. The shipbuilding budget is not entirely predictable, yet not purely stochastic. While planning and budgeting for particular programs begins years before the funds are actually appropriated and spent, costs may change unpredictably and Congress may make adjustments as it sees fit. Furthermore, laws governing appropriations require that major acquisitions such as aircraft carriers be funded out of a single year's budget, and this can result in

significant jumps in the shipbuilding budget from year to year. To complicate matters further, changes in world events, most notably wars, can lead to major swings in defense spending. As indicated by Gansler, historically, the defense budget has been quite volatile [28]. Consequently, this volatility should be considered in any model of the shipbuilding budget. To that, end we will assume that the defense budget also follows geometric Brownian motion. This leads us to the following model for the annual shipbuilding budget,

$$dB = gBdt + \sigma_B BdZ,$$

where g is the expected growth rate in the shipbuilding budget, and σ_B is the volatility of the budget. This model allows us to capture the general growth trend in the defense budget, but at the same time capture its volatility. Now there are some important caveats that must be mentioned regarding the use of geometric Brownian motion to model the budget. First, budgets are appropriated annually (though there is a limited ability to shift around funds within a budget year) while Brownian motion varies continuously. This would suggest that a stochastic process over a discrete time domain would be more appropriate. Admittedly, Brownian motion was largely selected for analytic convenience, but since we are considering the long-run, the relative impact of discrete time steps diminishes. Second, there are some cyclic features to the shipbuilding budget. For example if an aircraft carrier is procured every seven years, there will be a corresponding spike in the shipbuilding budget. However, as we look beyond the current budget planning cycle, this becomes less of a concern because there is a great deal of variability in the rate of acquisition since programs may be stretched out or accelerated to suit the particular needs of the time. So once again, the long-run view taken in this analysis diminishes the impact of this feature.

It was mentioned previously that defense production exhibits increasing returns to scale. To represent this behavior, we require a production function. A production function translates the rate of input consumption into a rate of output production,

in this case ships. A standard production function in economics is the Cobb-Douglas production function (see Varian, Chapter 1 [73]):

$$Y = AX^a.$$

The Cobb-Douglas production function facilitates the representation of economies of scale. When the parameter a is greater than 1, the production function exhibits increasing returns to scale. Constant returns to scale may be represented by setting $a = 1$ and decreasing returns to scale by setting $0 < a < 1$.

Since the entire shipbuilding budget is expended on building ships, the rate of input consumption is simply the ratio of the budget to the input cost.

$$X = \frac{B}{C}$$

Thus, we can define the output, $Y(t)$, as a function of input cost, $C(t)$, and the budget, $B(t)$.

$$Y = A \left(\frac{B}{C} \right)^a \tag{4.1}$$

Unfortunately, since B and C are governed by geometric Brownian motion, we cannot use Equation (4.1) as is. Instead, stochastic calculus is required to derive a model for $Y(t)$. Applying Ito's Lemma to Equation (4.1)(see Shreve, Chapter 4 [70]), we obtain the following stochastic differential equation for $Y(t)$

$$
\begin{aligned}
dY &= \left(ag - a\alpha_C + \frac{a(a-1)}{2}\sigma_B^2 - a^2\rho_{BC}\sigma_B\sigma_C + \frac{a(a+1)}{2}\sigma_C^2 \right) Y \, dt \\
&\quad + a\sigma_B Y \, dZ_B - a\sigma_C Y \, dZ_C
\end{aligned}
\tag{4.2}
$$

(see Appendix D for the derivation). Note that ρ_{BC} is the coefficient of correlation between the stochastic processes $B(t)$ and $C(t)$ such that $E[dZ_B dZ_C] = \rho_{BC}dt$. While it may not be obvious from Equation (4.2), the stochastic process $Y(t)$ is also governed by geometric Brownian motion with expected growth rate

$$\alpha_Y = ag - a\alpha_C + \frac{a(a-1)}{2}\sigma_B^2 - a^2\rho_{BC}\sigma_B\sigma_C + \frac{a(a+1)}{2}\sigma_C^2 \tag{4.3}$$

104

and volatility

$$\sigma_Y = a\sqrt{\sigma_B^2 - 2\rho_{BC}\sigma_B\sigma_C + \sigma_C^2}.$$

The proof of this is provided in Appendix D. Now, there is no reason to believe that the budget process, $B(t)$, is in any way correlated with the cost of ship input factors, $C(t)$. Thus, for much of this analysis it is assumed that $\rho_{BC} = 0$, but it is included here for completeness.

To determine the total number of ships produced over a given time interval, we must integrate $Y(t)$ over t. Let Y_T be the number of ships produced over the interval $[0, T]$:

$$Y_T = \int_0^T Y(t)dt.$$

To find the expected number of ships produced over the interval, we simply find the expected value of Y_T:

$$E[Y_T] = E\left[\int_0^T Y(t)dt\right] = \frac{Y_0}{\alpha_Y}(e^{\alpha_Y T} - 1).$$

Thus, we now have a basic model of the ship production process that describes the future output of navy ships, and this will serve as a basis for evaluating the efficacy of any potential acquisition process improvements.

One immediate consequence of the model is that if $\alpha_Y < 0$, then the rate of production is decreasing over time. This would occur if costs are increasing faster than the shipbuilding budget. Thus, we see that the model represents the phenomenon described earlier in the section where, in the absence of change, the Navy will eventually be unable to procure any ships at all. To make this more concrete, let us assume the notional parameter values listed in Table 4.2.

Note that we have initially set the the volatility parameters to zero. This constitutes the deterministic case and will serve to illustrate the basic impact of economies of scale. Note that costs are growing faster than the budget by 1% per year. However, if we examine the resulting output rate, we find that it is initially one ship per

Table 4.2: Notional Parameter Values

Parameter	Symbol	Value
Initial Annual Budget Rate	B_0	\$1 billion
Budget Growth Rate	g	2%
Budget Volatility	σ_B	0
Initial Input Unit Cost	C_0	\$1 billion
Expected Cost Growth Rate	α_C	3%
Cost Volatility	σ_C	0
Budget/Cost Correlation	ρ_{BC}	0
Cobb-Douglas Production Parameter	a	1.3
Cobb-Douglas Scaling Parameter	A	1

year, but that it is declining by 1.3% per year. As a result, only 17.6 ships will be built over the next 20 years. Note that since $a = 1.3$, the production function is exhibiting increasing returns to scale. Thus, even though the difference between the budget growth rate and the cost growth rate is 1%, the production level is declining at faster rate. This occurs because reductions in order quantity force the use of more inefficient production methods. The resulting waste means that there is an increase in the amount of input required to build a single ship. If ship production exhibited constant returns to scale ($a = 1$), the loss of production would mirror the budget shortfall at 1% per year. Consequently, 18.1 ships would be produced over 20 years as opposed to 17.6. This example clearly illustrates the dilemma faced by the ship acquisition enterprise. Cost growth actually accelerates the force structure shortfall when increasing returns to scale are present.

Of course, the reverse is true when we consider cost savings. Let us assume that the acquisition process improvements proposed by the Navy would instantaneously reduce input costs by 20%. Increasing returns to scale means that the ship production rate increases by 34%. Contrast that with constant returns to scale which would only boost the production rate by 25%. To put it more concretely, over the next 20 years, the Navy would be able to acquire 23.5 ships under increasing returns to scale compared to 22.7 under constant returns to scale.

Now what happens when there are fluctuations in the budget and input costs. An examination of Equation (4.3) reveals that the dependency of α_Y on the budget and cost growth rates, the budget and cost volatilities, and the production function exponent can be quite complex. We see that when there is no uncertainty in the input costs or budget, the growth rate behaves exactly as expected under increasing returns to scale. It is merely the difference between the budget and cost growth rates scaled by the production function exponent. When either the input costs or the budget is volatile, however, we see that the volatilities alter the production growth rate via quadratic functions of a. Budget uncertainty decreases the growth rate under decreasing returns to scale but increases it under increasing returns to scale. Cost uncertainty always increases the growth rate regardless of the economies of scale. Thus, we see that somewhat contrary to expectations, volatility in the input costs or budget can somewhat dampen the adverse impacts of increasing returns to scale when cost growth outpaces budget growth. Of course, this result should not be taken as absolute. Further examination of Equation (4.3) reveals that positive correlation between the cost and budget streams will rapidly erode the volatility induced boost in the production growth rate. Just to illustrate the point, if we increase the volatilities of both the cost and the budget streams to 0.02, then we find that the production rate is now only declining at 1.2% per year.

The model of shipbuilding developed in this section provides a means of evaluating the impact of cost savings or efficiency improvements on the production of Navy warships. But what is a change in production worth? To answer that question, the the price index approach described in Section 4.2 will be applied to the output of the acquisition model.

4.5 Valuing a Process Improvement to the Ship Acquisition Enterprise

The value of a process improvement is tied to the change in buying power that it entails. To assess the buying power both before and after a process improvement, we must a employ a price index. In the example problem presented in this analysis, we are concerned with the number of ships that the Navy may purchase. Therefore, the index must relate the price and number ships the Navy purchases with and without the process improvement in place.

First, we require the production rate under the current acquisition process. This will be designated as $Y_C(t)$. Next, we need the production rate after the process improvement is implemented, $Y_N(t)$. The price per ship is determined by simply dividing the budget rate by the production rate.

$$
\begin{aligned}
p_C &= \frac{B}{Y_C}, \\
p_N &= \frac{B}{Y_N}.
\end{aligned}
$$

Using the Laspeyres index, we get

$$
P_L = \frac{p_N Y_C}{p_C Y_C} = \frac{Y_C}{Y_N}.
$$

Since there is no substitution for ships with other military systems in our example, computing the other two indices discussed in Section 4.2 reveals that all are equivalent as expected:

$$
P_L = P_P = P_F = \frac{Y_C}{Y_N}.
$$

Since the implemented process improvement should increase the production level, the price index will be less than 1. Thus, ship prices have deflated, and the buying power of the shipbuilding budget has increased. To value this increase, we must translate it into monetary terms. This is accomplished by finding the amount of budget increase required under the current process to achieve a level of buying power equivalent to the reformed process. In other words, for a decision-maker to be indifferent

between the process improvement and the status quo, he would require an augmented budget stream. The size of the augmented budget rate is found by dividing the budget rate by the price index.

$$\frac{B}{P_F} = B\frac{Y_N}{Y_C}.$$

The monetary gain in buying power is determined by simply subtracting out the original budget rate.

$$B\frac{Y_N}{Y_C} - B.$$

To facilitate further analysis, let $G(t)$ represent the augmented budget stream.

$$G = B\frac{Y_N}{Y_C}.$$

Of course, this expression can be simplified given the definition of $Y(t)$.

$$G = B\left(\frac{C_C}{C_N}\right)^a$$

For analytic convenience, let us define a new process, $K(t)$, such that $K = C_C/C_N$. Applying Ito's Lemma yields

$$dK = (\alpha_{C_C} - \alpha_{C_N} - \rho_{CN}\sigma_{C_C}\sigma_{C_N} + \sigma^2_{C_N})K\,dt + \sigma_{C_C}K\,dZ_{C_C} - \sigma_{C_N}K\,dZ_{C_N}$$

where ρ_{CN} is the correlation coefficient between the current and new cost streams. Applying a logic similar to that presented in Appendix D, one can show that $K(t)$ is a geometric Brownian motion process with expected growth rate

$$\alpha_K = \alpha_{C_C} - \alpha_{C_N} - \rho_{CN}\sigma_{C_C}\sigma_{C_N} + \sigma^2_{C_N}$$

and volatility

$$\sigma_K = \sqrt{\sigma^2_{C_C} - 2\rho_{CN}\sigma_{C_C}\sigma_{C_N} + \sigma^2_{C_N}}.$$

$K(t)$ will prove useful later in the analysis, but for now it can be substituted back in to the definition of G to obtain

$$G = BK^a.$$

Applying Ito's Lemma to G results in

$$dG = \left(g + a\alpha_K + a\rho_{BK}\sigma_B\sigma_K + \frac{a(a-1)}{2}\sigma_K^2\right)Gdt$$
$$+\sigma_B GdZ_B + a\sigma_K GdZ_K. \qquad (4.4)$$

Of course, this too, is a geometric Brownian motion process with expected growth rate

$$\alpha_G = g + a\alpha_K + a\rho_{BK}\sigma_B\sigma_K + \frac{a(a-1)}{2}\sigma_K^2$$

and volatility

$$\sigma_G = \sqrt{\sigma_B^2 + 2a\rho_{BK}\sigma_B\sigma_K + a^2\sigma_K^2}.$$

If we again assume that the fluctuations in the budget are uncorrelated with the fluctuations in the cost of shipbuilding inputs (i.e., $\rho_{BK} = 0$), and we substitute for α_K and σ_K, we obtain

$$dG = \alpha_G Gdt + \sigma_G GdZ_G \qquad (4.5)$$

where

$$\alpha_G = g + a(\alpha_{C_C} - \alpha_{C_N}) + \frac{a}{2}(\sigma_{C_N}^2 - \sigma_{C_C}^2) + \frac{a^2}{2}(\sigma_{C_N}^2 - 2\rho_{CN}\sigma_{C_N}\sigma_{C_C} + \sigma_{C_C}^2) \quad (4.6)$$

and

$$\sigma_G = \sqrt{\sigma_B^2 + a^2\left(\sigma_{C_N}^2 - 2\rho_{CN}\sigma_{C_N}\sigma_{C_C}, +\sigma_{C_C}^2\right)}.$$

Thus, the model of $G(t)$ provides a means of evaluating the value of an acquisition process improvement provided that we can characterize the subsequent change in the cost stream. Consequently, within the context of our model, we must consider changes in acquisition costs in three ways: a change in the base cost level, a change in the cost growth rate, and a change in the cost volatility. Eventually, we will consider the impact of each of these changes, but first, to demonstrate the output of the valuation model, we will assume that there is a one time drop in the base cost level. Thus, the new cost structure is a fraction of the current [i.e., $C_N(t) = sC_C(t)$, where

$0 < s < 1$]. This means that the two cost streams have identical volatility and are perfectly correlated. In this case $G(t)$ simplifies dramatically to

$$dG = gGdt + \sigma_B GdZ_G.$$

If we let r be the risk-free discount rate, then the expected net present value of the increase in buying power when one switches from the current acquisition process to the improved acquisition process is

$$NPV = E\left[\int_0^\infty [G(t) - B(t)]e^{-rt}dt\right] = \frac{B_0}{r - g}\left[\left(\frac{1}{s}\right)^a - 1\right]. \tag{4.7}$$

when $r > g$. Note that the volatility of the budget and the cost streams has no impact on the expected value for this special case, and thus, it devolves into the form presented in Pennock, et. al., 2007 [57].

In order to make this result more concrete, it would be helpful to assign values to the model parameters. To that end, let us first consider the US Navy shipbuilding budget. Figure 4.2 depicts the US Navy shipbuilding budget from fiscal year (FY) 1980 to FY 2007 adjusted for inflation. The most obvious feature is that the budget is quite volatile. This is in part due to the previously mentioned legal requirement that the full procurement cost of a ship must be appropriated in a single budget year, but the impact of world events and political prerogatives are certainly evident as well. Notable features include the force structure buildup under the Reagan administration, as well as the drop in ship construction following the end of the Cold War. For the purposes of this analysis, we will focus on the post-Cold War trend in shipbuilding where the average inflation adjusted budget growth rate has been approximately 2.8% per year with a log-volatility of approximately 28%.

Since one of the assumptions of this model is that cost changes do not lead to substitution among military systems, we will focus on a single class of Navy ships, surface combatants. The primary surface combatant currently being procured by the Navy is the Arleigh Burke class guided missile destroyer (DDG-51). They are

US Navy Shipbuilding Budget
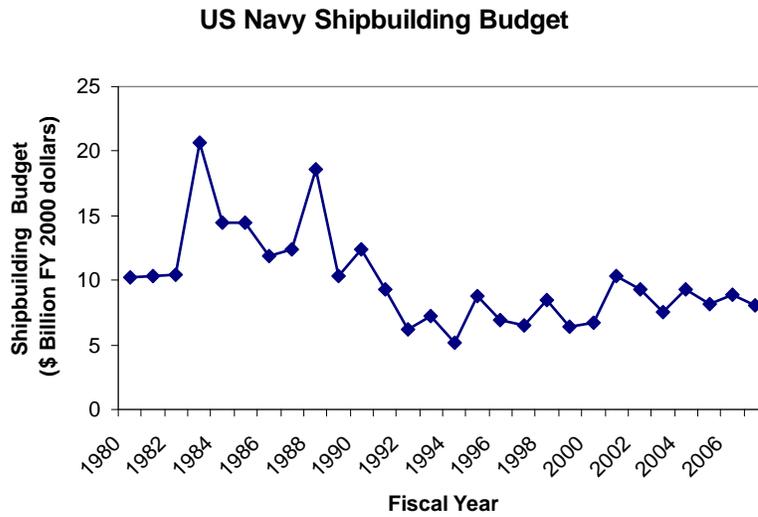
Figure 4.2: The US Navy shipbuilding budget over the period 1980 to 2007 in FY 2000 dollars

being commissioned at a rate of about 3 per year and cost about $1 billion apiece. Consequently, we will set the starting budget at $3 billion per year and the starting cost at $1 billion per unit. As far as cost growth, Arena and colleagues indicate that the cost of surface combatants has grown at a rate of 9.1% per year from 1965 to 2005 although approximately half of this cost growth is attributable to inflation [1]. Therefore, we will assume a cost growth rate of 4.5%. Unfortunately, Arena does not characterize cost volatility. As a proxy we will consider the defense price index published by the US Bureau of Economic Analysis. An analysis of this index over the period from 1947 to 2007 yields a log-volatility of approximately 3%. Since this index should in some sense capture the fluctuations in the cost of procuring defense systems, its volatility will serve as the proxy, albeit imperfect, for the cost volatility of ship procurement.

Finally, we must consider the production function. Unfortunately, production efficiencies in the manufacture of military systems are typically discussed in terms of

Table 4.3: Model Parameter Values

Parameter	Symbol	Value
Initial Annual Budget Rate	B_0	\$3 billion
Budget Growth Rate	g	2.8%
Budget Volatility	σ_B	0.28
Initial Current Cost Rate	C_{C0}	\$1 billion
Current Cost Growth Rate	α_{CC}	4.5%
Current Cost Volatility	σ_{CC}	0.03
Cobb-Douglas Production Parameter	a	1.3
Cobb-Douglas Scaling Parameter	A	0.719
Discount Rate	r	5%

learning curves rather than economies of scale, though the notion of efficient production rates is certainly recognized. Consequently, there is little data regarding the returns to scale in military shipbuilding. However, a GAO report regarding the F-22 program provides data correlating order size with unit cost [35]. When a Cobb-Douglas production function is fit to this data, the resulting production exponent is 1.35. To ere on the conservative side, a value of $a = 1.3$ will be used for the construction of surface combatants. Since the current production rate is 3 per year, the scaling coefficient A must be 0.719. Last but not least, we require a discount rate. The standard 5% discount rate will suffice. The assumed parameters are summarized in Table 4.3.

Returning to our model, if we apply these parameter values to Equation (4.7) and assume a 5% decrease in acquisition costs ($s = 0.95$), we find that the net present value of the increase in buying power is approximately \$9.4 billion. Thus, through this model, we can monetarily value a process improvement. But how does this compare to cost savings? If we set $a = 1$, meaning constant constant returns to scale, we obtain the NPV of the nominal cost savings. For this example the cost savings are \$7.2 billion, a 31% understatement of the true gain from implementing the process improvement. Clearly, using cost savings as the sole criterion for valuing acquisition process improvements can be quite misleading.

4.6 Valuing an Option to Improve the Ship Acquisition Enterprise

The previous section provided a method for valuing an improvement to the acquisition enterprise, but what if the efficacy of the improvement is uncertain? What if it turns out that the streamlining initiatives discussed in our example do not work in practice? In most reform efforts, there is a certain amount of technical risk. Technical risk is usually mitigated by staging the implementation of the new process. Each stage provides an exit point that allows decision-makers to terminate the effort if adverse information comes to light.

In our example, we will assume that there is a three stage process to implement a process improvement for ship acquisition[2]. The first stage is concept development and feasibility analysis. Since there are no actual acquisition programs involved in this stage, it should be relatively short and inexpensive. If the concept is determined to be infeasible or not cost effective, the process improvement project may be terminated at no additional cost. The second stage is a pilot test of the process improvement on the acquisition of a single ship. Failure in this stage will likely mean problems in the acquisition of the ship in question. In that event, we will assume that the Navy will still want to acquire the ship, and, consequently, rework will be required to complete the acquisition. Finally, the third stage is the enterprise-wide implementation of the acquisition process improvement. If there is a failure in this stage, substantial costs will be incurred because the acquisition of multiple ships will be adversely impacted. Since a low-risk, high-return project fares well under any decision criterion, the true value of the options approach lies in the evaluation of risky projects. Consequently, the values of the staging parameters were deliberately chosen to create a high risk of failure, and they are presented in Table 4.4.

---

[2]It should be noted that while the staging setup presented here is similar to that presented in Pennock, et. al., 2007 [57], the differences in the underlying shipbuilding model require a more sophisticated means of evaluation.

Table 4.4: Stage Parameter Values

Stage	Stage Cost ($ billions)	P(Success)	Rework Cost ($ billions)	Duration (years)
1	0.001	0.4	0	0.5
2	0.01	0.6	1	3
3	0.1	0.8	10	N/A

"Stage Cost" is the funding required to execute each stage, and "Rework Cost" is the cost incurred if the process improvement fails during the associated stage. "P(Success)" is the probability that the process improvement is successfully implemented in a given stage. Finally, "Duration" is the length of each stage. Note that the program costs are borne external to the shipbuilding budget.

The three stage implementation process essentially provides acquisition decision-makers with a series of options. At each stage they must decide whether or not it is worthwhile to continue the project. This is analogous to a compound call option, where buying the first option gives one the right to buy another option. What we would like to know is the value of this compound option. While the analytic valuation of compound options can be challenging, fortunately, straightforward numerical methods exist to evaluate complex options. Most are based on the principle that geometric Brownian motion can be approximated using a random walk, and one of the most popular is the binomial lattice method developed by Cox, Ross, and Rubinstein [15]. It functions by employing a random walk in which the state variable can only move discretely up or down. The moves are multiplicative, and the down move is the reciprocal of the up move. Thus, the resulting achievable state space forms a lattice, hence the name. With a discrete state space, the option is effectively a decision tree and can be solved using backwards induction. The binomial lattice method can achieve an arbitrary level of accuracy by reducing the size of the time step.

If we consider our previous example that valued a proportional drop in shipbuilding costs (See Equation (4.7)), we note that the net present value depended only on a single stochastic process, $B(t)$. However, when we took the expectation of the NPV, the volatility of $B(t)$ did not affect the value. Of course, that assumed there was no cost and no risk to achieve the process improvement. If, on the other hand, there are costs and risks, it would be prudent to implement the described three-stage process. This provides decision makers with the opportunity to exit the improvement effort prior to completion because of technical failures or unfavorable fluctuations in the shipbuilding budget. Therefore, we must consider the stochastic behavior of $B(t)$ in the valuation of the option. Fortunately, the binomial lattice method (or any other lattice method for that matter) can be used to account for this behavior. When we apply a lattice method to this option, we find that the net option value (NOV) of the option to implement the acquisition process improvement is approximately $1.09 billion. One may note that this is a considerable drop in value from the $9.4 billion calculated before we included the technical risk.

For comparison, the NPV of this improvement opportunity is approximately −$5.88 billion. Thus, NPV would imply that the Navy should expect to incur a loss if it initiates the acquisition reform project. The difference between the NPV and the NOV is attributable to NPV's failure to consider the value of staging. Thus, a decision-maker who utilized NPV as a decision criterion in the context of our example would incorrectly reject this reform opportunity.

Of course, it seems unlikely that the stochastic behavior of the cost stream would be exactly the same after the acquisition process improvement is implemented. In the following example, we will relax this assumption. Now, the new cost stream will have a higher volatility and only partial correlation with original cost stream. For the time being we will keep the cost growth rates identical in order to focus on the impact of the volatility. Unfortunately, relaxing these assumptions means that $G(t)$

Table 4.5: Market Risk Parameter Values

Parameter	Symbol	Value
Initial New Cost Rate	C_{N0}	\$0.95 billion
New Cost Growth Rate	α_{CN}	4.5%
New Cost Volatility	σ_{CN}	0.04
Cost Stream Correlation	ρ_{CN}	0.9
Budget/Cost Correlation	ρ_{BK}	0

is dependent upon the stochastic fluctuations in input costs. This means that its behavior is described by the stochastic differential equation (4.5), and the value of the process improvement is now dependent upon three stochastic state variables $B(t)$, $C_C(t)$, and $C_N(t)$.

Since the binomial lattice method can only accommodate a single stochastic state variable, another method is required to evaluate the process improvement option. A generalization of the binomial lattice method by Kamrad and Ritchken expands the lattice concept to handle an arbitrary number of stochastic state variables [44]. Like any dynamic programming method, however, Kamrad and Ritchken's method suffers from the curse of dimensionality. That is the state space becomes more and more unmanageable as the number of states increases. Thus, while we could certainly use the method to evaluate a three-state lattice, it would be more computationally efficient if we could reduce the dimensionality of the problem. It was observed earlier that the value of the process improvement depends on the ratio of the cost processes, and that this ratio is also governed by a geometric Brownian motion process, $K(t)$ (See Equation (4.4)). Therefore, we may evaluate the option to implement a process improvement with a two-state lattice over $B(t)$ and $K(t)$.

To evaluate the option where there is a volatility shift in the cost stream, we will use the previous parameter set (Tables 4.3 and 4.4), plus the additional parameters listed in Table 4.5 to model the market risk. There is still a drop a 5% drop in the base cost, but now each cost stream has its own volatility. Of course the cost streams will still be sensitive to many of the same economic perturbations, so the coefficient

of correlation was set to 90%. When we calculate the net option value, we find that it has jumped to \$2.08 billion. Why such a large increase? It is a well known result that volatility increases the value of an option. This is because of the downside risk mitigation provided by staging. Each stage provides decision-makers the opportunity to cancel the project due to an unexpected change in the input costs. Imagine if input costs declined precipitously. The savings realized through the increased efficiency of the acquisition process reform may no longer be sufficient to justify the costs to implement the reform. It should also be noted that since $G(t)$ is a function of the two cost streams, there is also an interaction effect at work. An examination of Equation (4.6) reveals that differences in the stochastic behavior of the old and new cost streams can also increase the growth rate of $G(t)$. This is evident in the NPV which has increased to $-\$4.9$ billion because of the difference in volatilities.

For the last example in this section, we will relax the assumption that the cost growth rate is unchanged after the acquisition process reform is implemented. In the end, the cost growth rate is really what matters. A drop in the base cost level of ship acquisition is certainly worth something, but as long as costs continue to grow faster than the shipbuilding budget, it is simply delaying the inevitable. Fortunately, even small decreases in the cost growth rate can have a major impact. If the cost growth rate drops below the budget growth rate, the shipbuilding enterprise becomes sustainable. To illustrate this point, let us reduce the input cost growth rate for the improved acquisition process a tenth of a percent (i.e., α_{CN} is reduced from 4.5% to 4.4%). This results in an NOV of \$3.95 billion, a substantial increase. If the process improvement were to achieve a cost growth rate of 3%, the net option value explodes to \$323 billion. Considering that the outlay for the first stage is only \$1 million, it is quite an attractive investment.

In order to better understand the general behavior of the model, sensitivity analysis was performed on the key model parameters, and the results are provided in

Figures 4.3 and 4.4. First, we consider the sensitivity of the NOV to the cost growth rates. Figure 4.3a reveals that as the growth rate of the new cost stream, α_{CN}, approaches the growth rate of the budget, the option value increases rapidly. This is because when the cost growth rate and the budget growth rate are the same, the production rate is sustainable. Similarly, as the budget growth rate, g, approaches the new cost growth rate, the same behavior is observed (Figure 4.3c).

Figure 4.3d reveals that changes in the two correlation coefficients yield opposing shifts in option value. As the correlation between the current and new cost streams, ρ_{CN}, increases, the value of the option decreases slightly. This is because when ρ_{CN} is high, the current and new cost streams tend to move together. Consequently, there are fewer opportunities to exploit favorable relative movements in cost. The reverse is true, however, for correlation between the budget and the cost ratio, $K(t)$. This occurs because the value of the process improvement is dependent upon the product of $B(t)$ and $K(t)$. Thus, correlated movements exaggerate the volatility of the value of the process improvement and create more exploitable opportunities. As mentioned previously, it is a well known result from options theory that an increase in volatility leads to an increase in the value of an option.

Finally, we consider the sensitivity of the NOV to the volatilities. Figure 4.4a reveals that increasing the volatilities of the current and new cost streams leads to opposite outcomes. This occurs because an increase in the volatility of the new cost stream has a positive impact on the growth rate of the value of the process improvement, α_G, while an increase in the volatility of the current cost stream has the opposite effect (See Equation (4.6)). If we consider an increase in the volatility of the budget, σ_B, Figure 4.4b reveals that it has a negligible impact under our assumption that the budget is uncorrelated with costs. If, however, we assume that the budget is correlated with the cost ratio ($\rho_{BK} = 0.5$), increasing budget volatility results in increasing NOV. The rationale is the same as for the impact of changing
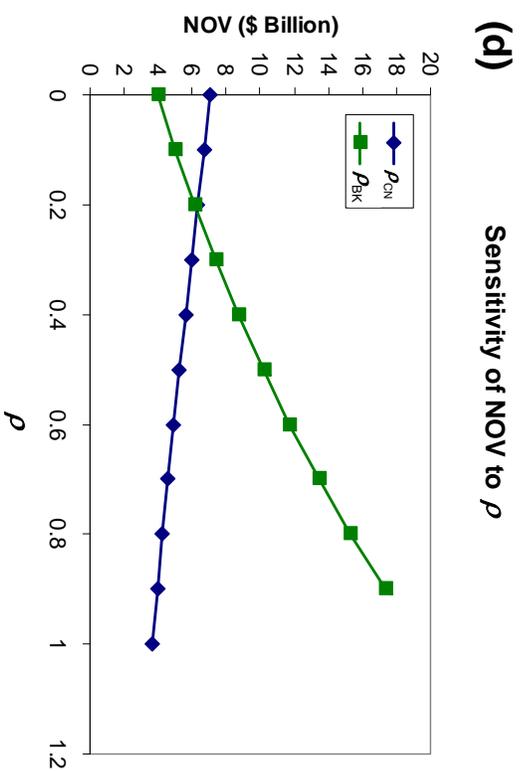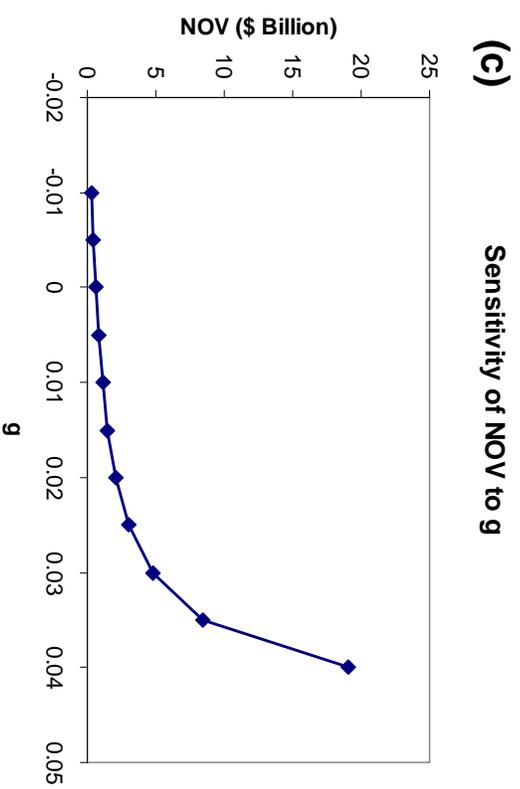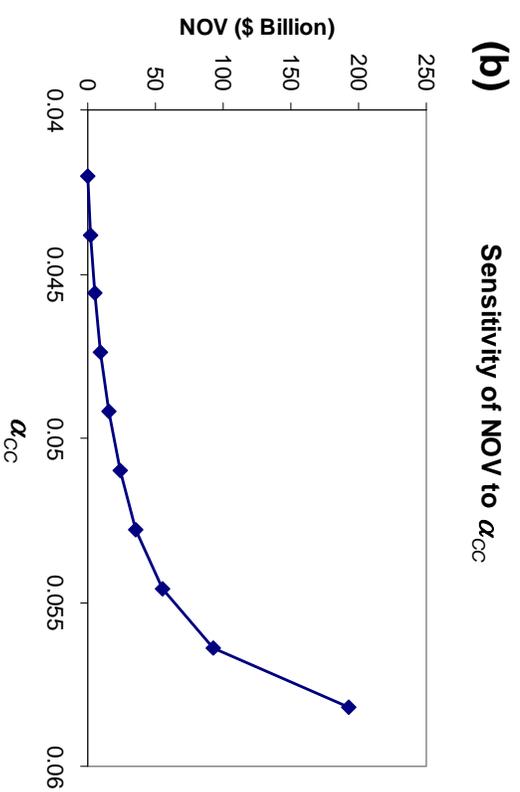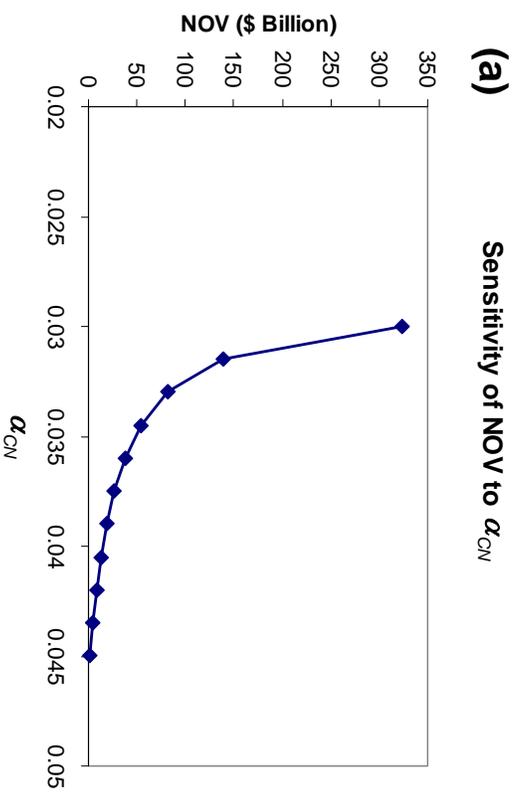
Figure 4.3: Model sensitivity results for α_{CN}, α_{CC}, g, ρ_{CN}, and ρ_{BK}

ρ_{BK}.

Through the series of examples presented above, it has been demonstrated that it is possible to value acquisition process improvements monetarily while still considering the fact that efficiency improvements change what is bought. Furthermore, the application of the options approach reveals that failure to consider the risk mitigation inherent in staged investments can cause decision-makers to reject otherwise valuable opportunities. It is important to note that all of the examples thus far have assumed risk-neutrality on the part of the decision-maker. While this should theoretically be true in a government context, in practice, most individual decision-makers are risk-averse. Since it would be impractical and possibly detrimental to include risk aversion in the option valuation, the next section will present a portfolio method for evaluating possible acquisition process improvements that will allow decision-makers to trade-off risk and return.

4.7   A Portfolio Approach to Investment Selection

The portfolio approach accounts for risk attitudes by evaluating an investment in terms of both the expected return as well as the uncertainty in the return[3]. By considering the whole portfolio of possible investment options, decisionmakers are able to explicitly trade-off risk and return. Return in this context would be the net option value calculated using the method developed in this chapter. As for risk, there are many possible metrics, but here we consider two, the probability of a loss and the conditional expected loss. The probability of a loss is the likelihood that any loss at all occurs, and the conditional expected loss is the expected loss assuming that one occurs. These are both characterizations of downside risk. The portfolio concept is illustrated in Figure 4.5. Alternative investments, denoted by the Ps, are plotted in

---

[3]It should be noted that the portfolio concept for acquisition improvements was first introduced in Rouse, et. al., 2006 [65] and was subsequently expanded in Pennock, et. al., 2007 [57]. This section represents a further evolution of that concept.
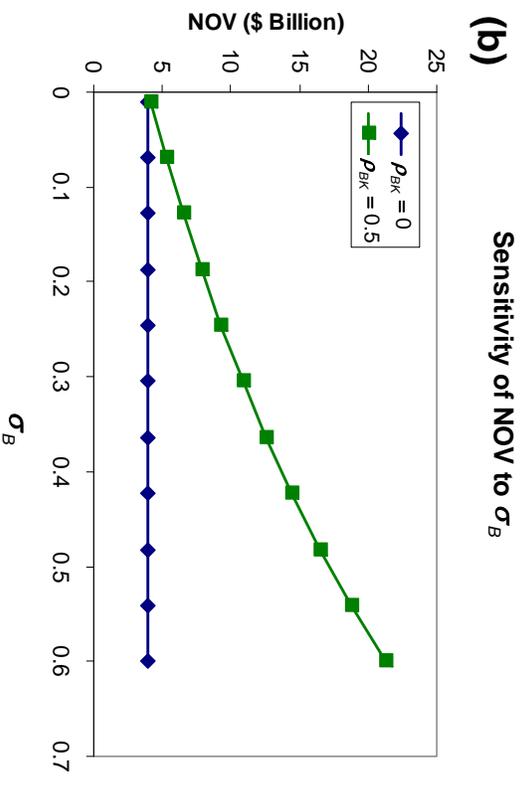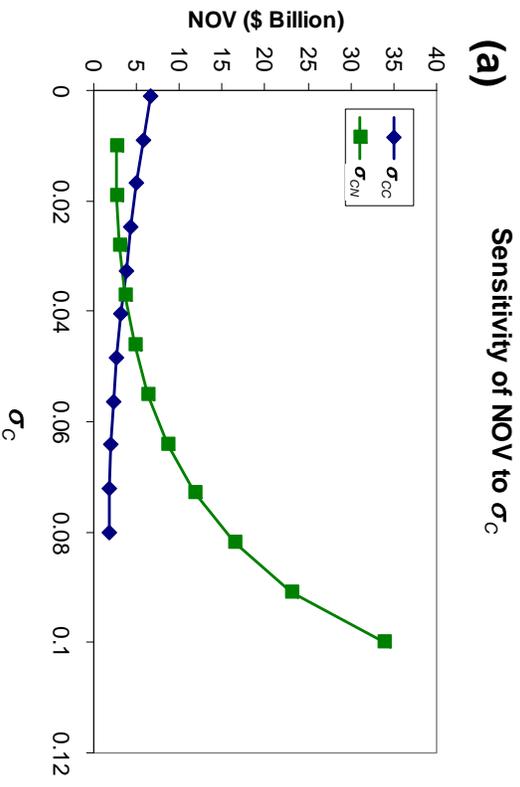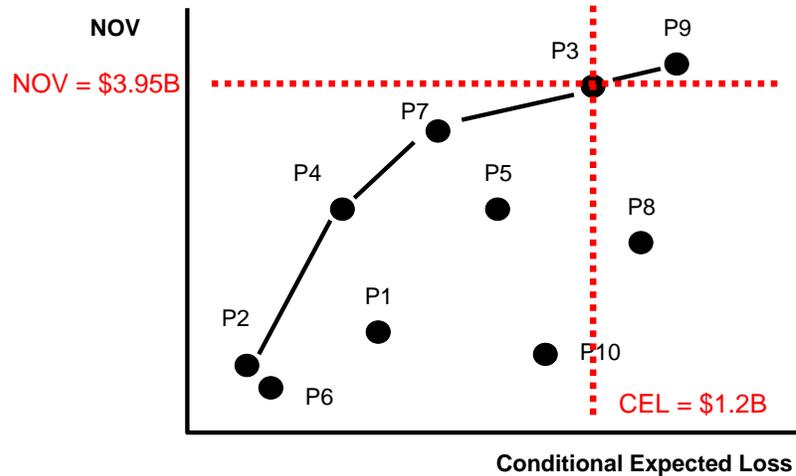
Figure 4.4: Model sensitivity results for σ_{CN}, σ_{CC}, σ_B

terms of expected return and risk. In this case, the risk is expressed as the conditional expected loss.

Figure 4.5: Notional portfolio of acquisition reform projects.

One may note that this is conceptually somewhat similar to Markowitz portfolio theory [49–52]. Under Markowitz portfolio theory, securities investments are also plotted based on their risk and return, and individual securities may be combined in portfolios to create new investment options with different risk and return characteristics. The risks and returns of these new portfolios are simply linear functions of the risks and returns of the constituent securities. The non-dominated set of portfolios constitutes the efficient frontier, and an investor would only ever purchase a portfolio from this frontier.

Where the analogy breaks down is in the combination of process improvements. More than likely, the simultaneous implementation of multiple process improvements will lead to some interaction effects. If two process improvements were to have a synergistic effect, then the return from combination of the two would be greater than the sum of the individual returns. Alternatively, the two process improvements could interfere with each other and actually reduce the realized benefit from implementation.

Thus, the evaluation of combinations of process improvements is not as straightforward as in Markowitz portfolio theory. Even so, it is still possible to consider a set of process improvements, but it would require modeling the set altogether and reevaluating the risk and return. This would essentially constitute a new project that would be placed in the portfolio plot.

Returning to the portfolio plot, we can see that the set of acquisition process improvements has an efficient frontier just as in Markowitz theory. This is indicated by the solid lines in Figure 4.5. A decision maker evaluating acquisition process improvements for possible investment would only want to choose one from this frontier. Take, for instance, P4 and P5. Both provide the same return, but clearly P5 is more risky. Consideration of the efficient frontier allows a decision-maker to find a project that provides the appropriate balance of risk versus return.

In order to understand the motivation for the portfolio approach to acquisition reform investments, it is necessary to provide some discussion of risk attitudes. The previously described option valuation method implies that the decision-maker is risk-neutral. A risk-neutral decision-maker is one who is indifferent between the expected value of a risky return and the equivalent lump sum. In contrast, a risk-averse decision-maker would prefer the certainty of the lump sum. Generally speaking, one would expect a large organization with sufficient resources to safely absorb any potential loss to be risk-neutral. However, the decision-makers within the government are not likely to view a failed initiative as very favorable for their careers, and thus, they would probably exhibit risk aversion. The standard way to handle risk aversion in real options analysis is through market mechanisms, but as explained earlier, this is not feasible in a government context. Another way to handle risk attitudes is through multi-attribute utility theory. Utility theory involves assessing a decision-maker's preferences quantitatively and translating them into a dimensionless measure through utility functions. This is impractical in this context for two reasons. First,

since these decisions are being made in the public domain, they are subject to scrutiny by elected officials and the public. A decision-maker's personal utility score will not likely satisfy either group as sufficient justification for a decision with implications for the well-being of US national security. Second, utility would effectively eliminate the monetary valuation we just developed and would complicate assessment of any return on investment.

The portfolio approach provides a compromise position. Projects to improve acquisition processes are still assessed using a risk-neutral approach, but we may also assess the probability distribution of their outcomes. This will allow us to extract measures of uncertainty for consideration by decision-makers. There is a latent inconsistency in this approach, however. There is an implicit assumption that after the initial investment decision, the decision-maker will behave in a risk-neutral manner for the subsequent stages. With that caveat, it is fairly safe to say that most decision-makers would likely prefer some description of the risk, albeit imperfect, to none at all.

Assessing the distribution analytically would be difficult if not impossible in most practical situations. Instead, Monte Carlo simulation may be used to approximate the distribution of an option's value. Since the lattice method employed previously is essentially a random walk combined with a decision tree, it is trivial to generate sample paths over the lattice. The value received when the final stage of the project is implemented is more challenging. If there is no analytic solution for the distribution of the terminal value, it must be simulated as well. For the example problem presented in Section 4.7, we must simulate the Brownian motion paths of the gain in buying power. Once again, a random walk is the most straightforward means to approximate Brownian motion, and this can be accomplished by simply extending the lattice. Since the extended lattice must still have a finite horizon, the expected present value of the remaining cash flow serves as a terminal value. Thus, the longer the lattice horizon,

the more accurately variation will be captured.

To illustrate this method, we will return to our acquisition process improvement example from Section 4.7. In particular, we will consider the case were $\alpha_{CN} = 4.4\%$ and $\alpha_{CC} = 4.5\%$. The Monte Carlo simulation was run for 1 million iterations with an extended lattice horizon of 100 years. Of course, the expected value of the resulting distribution is the same as the NOV calculated in the last example of Section ($3.95 billion). The variation in the NOV is substantial, however, with a standard deviation of $134 billion. Why such a large spread? First, the discrete nature of the staging combined with significant technical risk means that there is a mode corresponding to the losses incurred from failure at each stage. Second, if the process improvement is successfully implemented, the high volatility of the budget means that there is a large spread in the realized increase in buying power. Consequently, we are faced with a fairly complicated probability distribution that would be quite difficult for decision-makers to interpret. Thus, the motivation for the risk measures described earlier is apparent.

Returning to the example, the probability of a loss for this acquisition process improvement is approximately 83%, and the conditional expected loss is approximately $1.2 billion. We may use the net option value in conjunction with the conditional expected loss to place the project in the portfolio depicted in Figure 4.5. Furthermore, the two risk measures also tell us that there is a very high probability that this project will fail, but if it does fail, the expected cost is relatively low compared with the potential returns.

While the portfolio approach is not a perfect means to consider risk, it does provide decision-makers with a method for harnessing a monetary valuation of acquisition reform initiatives in conjunction with a justifiable means of risk valuation. This will lend significantly more credibility to acquisition reform decisions. It is much more likely that elected officials will accept a high probability of loss as a reason for

dismissing an otherwise high-value project than a low utility score.

4.8 Summary and Policy Implications

The lack of objective criteria for evaluating defense outcomes has long been a stumbling block in defense policy debates. When we consider the domain of defense acquisition reform, the problem is somewhat simplified. It is not necessary to determine what to buy, only how to buy it more efficiently. To that end, the traditional approach to evaluating process improvements is to consider nominal cost savings. The analysis presented in this chapter, however, has demonstrated that cost savings can significantly understate the value of an acquisition process improvement. Economies of scale within the acquisition enterprise induce a non-linear response to cost reductions. Thus, the government does not buy the same things for less, it changes what and how much it buys. To account for the change in buying power that results from an acquisition process improvement, price indices similar to those used to track inflation were introduced. Price indices allow for the determination of an augmented budget that is equivalent to the increase in buying power. This augmented budget provides the basis for valuing an acquisition process improvement monetarily.

Of course, there is always risk involved in any process change. Market risk and technical risk can reduce the value of an investment opportunity. Fortunately, staging an investment can mitigate downside loss and restore some value. The traditional means of investment valuation, NPV, fails to account for the staging effect and, thus, can significantly under value investments. Options analysis, on the other hand, overcomes the shortcomings of NPV and appropriately values staged investments. Coupling the monetary valuation of buying power with options analysis yields a valuation method that remedies the shortcomings of past approaches.

Even with staging, some risk remains. Thus, it is imperative to provide decision-makers with a means to trade between risk and return. The portfolio method developed in this chapter does just that. By plotting potential acquisition reform projects by their return and a risk metric, decision-makers may compare alternatives and find one that presents the best balance of risk and return. Most importantly, the basis of the decision is justifiable, a crucial characteristic in government policymaking.

The method presented here not only makes strides to remediate the lack of objective decision criteria to support defense acquisition policymaking, but it also leads to some policy implications. The first is that even small changes can have tremendous value. Such opportunities may have been overlooked in the past due to the shortcomings of traditional valuation methods or the lack of any valuation method at all, but when economies of scale and the risk mitigation effects of staging are properly considered, their value becomes apparent. Second, to achieve a sustainable acquisition program, it is necessary to look beyond just production. Concepts such as six sigma and lean manufacturing are certainly beneficial, but they will not solve the affordability problem of military systems. Instead, one must consider the entire enterprise from defense authorizations in Congress to the management of the supplier base. Thus, this chapter echoes the theme presented in Chapter 3. Only by treating the entire acquisition enterprise as a system will it be possible to find effective solutions to the problems that have plagued defense acquisition for over 50 years.

CHAPTER V

CONCLUSIONS AND FUTURE RESEARCH

Despite over 50 years of effort and hundreds, if not thousands, of additions and alterations to the laws and regulations governing defense acquisition, most concerned participants are not satisfied with the operation of the defense acquisition enterprise. There is a general sense that military systems take too long and cost too much to develop and acquire. While certainly some of this disappointment can be attributed to unrealistic expectations, nearly every single defense acquisition program finishes over budget and behind schedule. Numerous studies and audits have revealed that the normal operation of the acquisition system leads to waste and delay. Among the causes are the extensive use of immature technology, significant concurrency between program phases, and unstable funding and requirements.

In Chapter 1, three possible contributors to the failure of acquisition reform were identified: misalignment of incentives, a lack of systems view, and a lack of objective evaluation criteria. This dissertation considered each of these three factors in turn. In Chapter 2, the acquisition system was modeled as a game where stakeholders in military acquisition competed to meet their respective objectives for deployed capability. The game revealed that a tragedy of the commons is at work where participants in the defense acquisition system are incentivized to act in contradiction to regulations.

In Chapter 3, a simulation of the defense acquisition system was developed to evaluate the latest acquisition transformation initiative, evolutionary acquisition. The simulation revealed that proponents of evolutionary acquisition have been overly focused on individual programs rather than the system as a whole. While evolutionary

acquisition should reduce the cost and increase the performance of individual acquisition programs, it could potentially raise the cost of operating the acquisition system as a whole.

Finally, in Chapter 4, an assessment method was developed to monetarily value and compare improvements to the defense acquisition system. A comparison of the developed method to traditional means of evaluation revealed that failure to consider the full set of economic forces at work can lead to a significant underestimation of the impact of a process improvement. It also demonstrated that it is possible to quantitatively asses the impact of acquisition reform.

These results indicate that there is some validity to the suggested causes of acquisition transformation failure. There is an incentive for stakeholders in the defense acquisition system to push for technology that is more immature than recommended by DoD best practices, and clearly, the lack of a systems view has lead to unrealistic expectations from evolutionary acquisition policies. Finally, the traditional means of assessing cost savings tends to undervalue improvements to defense acquisition and potentially leads to erroneous assumptions regarding the validity of process improvements.

5.1 Recommendations for Improving the Efficacy of Defense Acquisition Transformation

A contemplation of the results presented in the previous chapters naturally leads to several recommendations to improve the efficacy of efforts to transform the defense acquisition system. While it is not expected that the implementation of these recommendations will cure all that ails defense acquisition, the sheer size and import of the enterprise means that even small improvements can have significant payoffs.

First, before any change in acquisition policy is instituted, policymakers should consider whether the incentives are aligned with the objectives of the policy. Often,

acquisition policy changes are declarations of what should be without any consideration of the context in which the declaration is imposed. One must assume that all participants will act in their own best interests. This is not to suggest that all participants in the defense acquisition system are selfish. Rather, different stakeholders will have different views on the priority and import of the various outputs of the acquisition system. For example, an Air Force officer may perceive that global precision strike from air or space assets is the key to realizing US national security goals. A Navy officer, on the other hand, may feel that the access and power projection provided by sea basing is the imperative. Consequently, both will seek to pursue what they believe is in the best interests of the United States, possibly in opposition to each other. The result is that participants in the defense acquisition enterprise may harm the realization of their own goals even when behaving rationally. Such outcomes constitute a classic tragedy of the commons where acquisition programs serve as common resources to be exploited by multiple stakeholders. The result is over-exploitation of the resources to the detriment of all.

Consequently, this type of behavior may reduce the efficacy of acquisition reform policies. There are really only two potential solutions to this problem. The first is oversight and enforcement. Compliance with rules and regulations would be closely monitored and strictly enforced. Unfortunately, this tends to lead to a significant amount of overhead, and participants may still find ways to "game" the system. Second, design the policy such that participants are incentivized to comply. In other words, with the proper incentives, actors in the defense acquisition system will behave in the desired manner even without enforcement. In the commercial world this can be achieved by establishing ownership or tying compensation to the desired outcome. The second option is really preferable, but it may be challenging to craft incentives in a public sector environment. In reality, some combination of the two approaches will likely be required.

The next major recommendation is that potential acquisition policies should be analyzed in the context of the entire defense acquisition system before they are implemented. Chapter 1 discussed the vacillation between extremes in defense acquisition reform initiatives (e.g., increased oversight versus streamlining and COTS versus MIL-SPEC). The defense acquisition enterprise is a large and complex system. Any change to that system is likely to entail unintended consequences, and consequently, there are tradeoffs to consider regarding the imposition of any acquisition policy. The vacillation in acquisition policy is in large part attributable to a failure by policy makers to consider the tradeoffs inherent in their decisions. Only through systems modeling and analysis is there any hope of understanding the extent of a policy's impact.

An example of the application of systems modeling was presented in Chapter 3 where evolutionary acquisition technology policies were examined in the context of the entire defense acquisition system. The analysis revealed that what was expected to be an all-around improvement in terms of the cost, performance, and speed of acquisition programs through the use of evolutionary policies is actually a tradeoff. One can use evolutionary acquisition to improve performance and speed but possibly at additional cost. Similar analyses should be performed on every acquisition reform considered by the DoD, Congress, or the President.

Finally, the third major recommendation is that policymakers should systematically consider the full economic impacts of their decisions. The DoD is in the unique position that it is involved in nearly every piece of the defense value chain and is virtually the only customer. As a result, it cannot share the burden of maintaining the defense industry with others. Any decision that it makes affects the health and efficiency of the defense industry and can have wide-ranging implications. For example, a decision to cancel a program could cause a supplier of key components to go bankrupt. The next time those components are required, the DoD must essentially rebuild that competency at considerable cost in time and money. This is

not to suggest that these factors are not considered by policymakers. They certainly are. Rather, a lack of appropriate analysis tools has hindered their ability to fully understand the impacts, and the tools developed in this dissertation constitute at a step toward remediating that problem. A more systematic approach to assessing value using appropriate economic and investment tools could lead to more objective and interpretable measurements of outcomes. Objective measurement of outcomes means that policy alternatives can be compared and contrasted as was demonstrated in Chapter 4.

5.2 Recommendations Regarding the Implementation of Evolutionary Acquisition

Since a significant portion of this dissertation dealt with evolutionary acquisition, there are recommendations specific to that initiative that logically follow from the findings presented in this dissertation. They are as follows:

First, technology maturity requirements should not be optional. The intent of evolutionary acquisition is to create shorter acquisition cycles that make more modest, or evolutionary, increases in system capability. Immature technology introduces significant cost and schedule risks that preclude short acquisition cycles. Chapter 2 revealed that for systems that provide multiple capabilities, stakeholders are incentivized to push for immature technology. Thus, the DoD should not expect compliance with evolutionary acquisition technology recommendations when either maturity requirements are optional or exemptions are routinely granted.

Second, evolutionary acquisition is more appealing for programs with low development and procurement costs. While evolutionary acquisition has the potential to increase the average performance of systems in the field, this performance gain may come at additional cost. Shorter acquisition cycles mean that development and procurement costs are incurred more frequently than under traditional acquisition methods. Of course, the lower these costs are, the more cost effective evolutionary

acquisition becomes. Additional research is required to definitively understand how shifting to evolutionary acquisition affects these cycle costs.

Third, the costs of evolutionary acquisition can be controlled by increasing the time between acquisition cycles. If deploying new systems at the fastest possible rate is too expensive, simply insert a delay between cycles. This results in some sacrifice of average performance in the field, but allows for a stable acquisition technology policy while still managing costs. A possible trade-off, however, is that periods of inactivity may adversely impact the industrial base.

Fourth, evolutionary acquisition is more important for military specific technologies. Technology development is dependent upon experimentation, testing, and learning. For technologies that are unique to the military, the only sources of information on the performance and shortcomings of technologies is through the deployment and use of military systems. Slower acquisition cycles mean slower learning. Consequently, evolutionary acquisition, through its faster acquisition cycles, has the potential to increase the rate of improvement of military technologies.

Fifth, evolutionary acquisition depends on a well-funded R&D system. Since acquisition programs require mature technology under evolutionary acquisition, it is incumbent upon the defense R&D system to oversee the timely maturation of technologies. If this process is starved of funding, the DoD would be better off employing traditional acquisition approaches.

Finally, as the speed of acquisition cycles increase, the rate of technology development becomes the limiting factor. When acquisition cycles proceed faster than the development of new technologies, there is no gain to deploying a new system. Thus, if evolutionary acquisition successfully reduces acquisition cycle time, the full benefits may not be realized if the technology development process serves as a bottleneck. Therefore, a well-managed technology development process that aligns technologies

and funding with needs is imperative to successfully implementing evolutionary acquisition.

5.3 Future Work

The results presented in this dissertation suggests several avenues for future work. As stated in Chapter 1, there is a decided lack of academic research in the field of defense acquisition, and consequently, there is a plethora of research questions that merit a substantial amount of attention from researchers.

The first and perhaps most pressing need, is the application of game theory to additional areas in the defense acquisition enterprise. There have been some instances of its application in the past, but not nearly to the extent merited by the magnitude of the problem. Chapter 2 revealed a tragedy of the commons at work regarding technology policy, but there are many other problem areas that merit further study. Among these are competition both between programs and between the services for funding, the behavior of program managers, and the bidding by and competition between defense contractors. A closer examination of these topics will likely reveal the underlying mechanisms that drive the undesirable behavior often exhibited by participants in defense acquisition. Understanding these mechanisms is a prerequisite to designing effective policies to combat such behavior.

The importance of this research topic is evident in the history of defense acquisition. Countless policy reforms have resulted in little or no change in the behavior of participants in the defense acquisition enterprise. No policy, no matter how well conceived, can be effective when it is ignored or circumvented. Consequently, a better understanding of the motives of participants in defense acquisition is imperative to successful transformation.

Second, the analysis of evolutionary acquisition presented in Chapter 3 is really

a first step in analyzing a complex problem. The simulation model of defense acquisition, in particular, could be extended to answer additional questions regarding evolutionary acquisition and the operation of the defense acquisition system. First and foremost, further investigation into the relationship between the targeted level of capability improvement of an acquisition cycle and the cost and duration of system development, production, and deployment is required. Understanding this relationship is critical to fully understanding the cost implications of evolutionary acquisition. Additionally, the simulation could be modified to consider concurrency among program phases, acceleration of technology development, operations and maintenance costs, system integration issues, acquisition program budgeting issues, alternative technology management schemes, and the inclusion of the JCIDS and PPBE. While it is unlikely that the inclusion of any of these factors will alter the basic findings of Chapter 3, a better understanding of these issues will lead to a more detailed understanding of the tradeoffs inherent to evolutionary acquisition. Furthermore, understanding how the pieces of the defense acquisition enterprise interact would hopefully lead to more effective acquisition policies on the whole.

Finally, the process improvement valuation method developed in Chapter 4 is a proof-of-concept. Since it is fairly theoretical, substantial additional work would be required to transform it into a practical tool for decision-makers. Among the required improvements would be the addition of discrete state variables where appropriate, more complex production functions, and consideration of force structure reallocations. The monetary valuation of an improvement program significantly enhances the ability to justify its implementation. The value of such justification in public policy cannot be overstated. Quantitative assessments of value significantly improve the ability of policy makers to objectively compare and contrast policy alternatives as well as garner support for their implementation.

5.4   Summary

The defense acquisition enterprise is a unique institution, and it is often criticized for its inefficiencies. However, citizens, elected officials, and even members of the defense department itself must realize that defense acquisition will never be as efficient as the private sector. The economics of the situation simply preclude it. The government essentially funds and manages an entire industry by itself. In order to retain certain capabilities that have no other commercial application, the government must pay a premium to preserve the required expertise and equipment even when they are not in use. Furthermore, the public nature of defense funding precludes the efficient budgeting and management of acquisition programs. Even so, there are certainly ways for the Department of Defense to get more for its money within the constraints in which it operates.

Acquisition reform has been a popular issue with politicians for over 50 years, yet repeated attempts to transform the acquisition enterprise have yielded little benefit. In fact, reform initiatives often impose the same policy changes as previous attempts without any consideration of why the previous effort failed. The logical question then is why have so many transformation efforts failed? Three contributing factors to transformation failure were identified and analyzed in this dissertation, and recommendations were made to remediate these issues. It is expected that these recommendations, if implemented, may help, but they are not likely to produce drastic changes in the cost of military systems. The drivers of system cost increases involve far more than just inefficiency. Ultimately, there are no easy answers, but the importance of the defense acquisition enterprise and the proper stewardship of public funds are sufficient justification to merit the continued expenditure of effort to achieve the best possible performance.

APPENDIX A

SUPPLEMENT TO CHAPTER 2

This appendix contains the proofs of the theorems presented in Chapter 2. The following lemma is required for the proofs and so is presented first.

Lemma A.1. *Given the assumptions of the model and symmetric technologies, $Y(g,n) = A(n)W(g) + I$ where $1 \leq A(n) \leq n$ and $A(n)$ is strictly increasing in n.*

Proof. First, we will prove that $Y(g,n) = A(n)W(g) + I$. We know from the model assumptions that $Y(g,n) = E[X_M] + I$, so we must show that $E[X_M] = A(n)W(g)$. Recall that $X_M = \max(X_1, X_2, \ldots, X_n)$, and that the distribution of X_i will vary with g (i.e., $X_i \sim F_X(x; g)$). As indicated in Section 2.2, given the distribution of X_i for any one value of g, we can determine the distribution of X_i for every other value of g through an affine transformation on X_i. Thus, we only need to specify one distribution to obtain the rest. To that end, let g_1 be the value of g such that $W(g_1) = 1$. Let us assume that we know distribution of X_i for the case $g = g_1$ (i.e., $F_X(x; g_1)$). For notational purposes let us designate X_i for this particular of value of g_1 as D with the the distribution of $F_D(d) = F_X(x; g_1)$. Consequently, $E[D] = 1$, and the value of X_i for any value of g can be found through the affine transformation $X_i = W(g)D$. Thus, we can state that

$$
\begin{aligned}
X_M &= \max(W(g)D_1, W(g)D_2, \ldots, W(g)D_n), \\
X_M &= W(g)\max(D_1, D_2, \ldots, D_n), \\
X_M &= W(g)D_M, \\
E[X_M] &= W(g)E[D_M].
\end{aligned}
$$

Since $E[D_M]$ is a function of n but not g, we can state that $E[X_M] = A(n)W(g)$. It follows directly that since $E[D_M]$ is strictly increasing in n that $A(n)$ must be as well.

Finally, we prove that $1 \leq A(n) \leq n$. From the first part we know that

$$A(n) = \frac{E[X_M]}{E[X]} = n \frac{\int_0^\infty x F^{n-1}(x; g) f(x; g) dx}{\int_0^\infty x f(x; g) dx}.$$

Since $E[X_M]$ is greater than $E[X]$, $A(n) \geq 1$. Since $0 \leq F^{n-1}(x; g) \leq 1$ the integrand of the numerator must be less than the integrand of the denominator. Therefore, the fraction must be less than 1 and $A(n) \leq n$ □

Proof of Theorem 2.1. This proof consists of two parts. First, we must prove the existence of a unique optimum. Second, we must prove that the optimal solution decreases as n increases.

Since the policy is symmetric, we can assume that $g = g_1 = g_2 = \cdots = g_n$. Thus, the optimization problem reduces to

$$\max_g \quad (1+g)^{V(g,n)} - 1.$$

Without impacting the optimal solution we can transform the objective function by dropping the -1 and taking the natural log.

$$\max_g \quad V(g, n) \ln(1+g)$$

Applying the first order optimality condition we obtain

$$\frac{\partial V(g, n)}{\partial g} \ln(1+g) + \frac{V(g, n)}{1+g} = 0.$$

Substituting for $V(g, n)$ yields

$$-Y'(g, n) \frac{\ln(1+g)}{Y^2(g, n)} + \frac{1}{Y(g, n)(1+g)} = 0.$$

Rearranging terms we obtain

$$Y(g, n) = Y'(g, n) \ln(1+g)(1+g). \tag{A.1}$$

139

Using Lemma A.1 we know that $Y(g, n) = A(n)W(g) + I$. Substituting into Equation (A.1) yields

$$A(n)W(g) + I = A(n)W'(g) \ln(1 + g)(1 + g).$$

If we rearrange terms we obtain

$$W'(g) \ln(1 + g)(1 + g) - W(g) = \frac{I}{A(n)}. \qquad (A.2)$$

Taking the first derivative of the left hand side yields

$$W''(g) \ln(1 + g)(1 + g) + W'(g) \ln(1 + g) > 0.$$

Given the modeling assumptions made previously, the left hand side is strictly increasing over the domain $g \geq 0$. Since the right hand side is a constant, then there exists at most one solution g^* over the domain $g > 0$. If the first order condition is not satisfied over the domain, then the optimal solution is $g^* = 0$. To prove that $g^* > 0$ achieves the maximum we must apply the second order condition.

$$-\frac{Y''(g, n)}{Y^2(g, n)} \ln(1 + g)(1 + g) + 2\frac{(Y'(g, n))^2}{Y^3(g, n)} \ln(1 + g)(1 + g) - 2\frac{Y'(g, n)}{Y^2(g, n)} - \frac{1}{Y(g, n)(1 + g)}$$

Substituting the solution of the first order condition into the second order condition yields

$$-\frac{Y''(g^*, n)}{Y'(g^*, n)Y(g^*, n)} - \frac{1}{Y(g^*, n)(1 + g^*)} < 0.$$

for $g^* \geq 0$. Thus, we see that g^* is the optimal solution. This proves the first part.

To prove the second part, we return to the first order condition found in Equation (A.2). From Lemma A.1 we know that $A(n)$ increases with n. Therefore, the right hand side decreases with n. Since the left hand side is strictly increasing in g but constant with respect to n, then we know that g^* decreases with n. Since the long-run arrival rate, $1/Y(g, n)$, of the renewal process decreases as n increases for all values of g, then the long-run effective growth rate must also decrease for all values of g. Therefore, the effective growth rate for $g^*(n+1)$ must be less than that for $g^*(n)$. $\quad \square$

Proof of Corollary 2.1. Using the first order optimality condition from the proof of Theorem 2.1

$$W'(g)\ln(1+g)(1+g) - W(g) = \frac{I}{A(n)},$$

we see that as I increases the right hand side increases. Since the left hand side is strictly increasing in g, g^* must also increase. Since the arrival rate decreases for all g when I increases, the long-run effective annual growth rate decreases as I increases. □

Proof of Theorem 2.2. The best response for any given stakeholder is

$$\max_{g_i} (1 + g_i)^{V(\boldsymbol{g},n)} - 1.$$

Following arguments similar to the proof of Theorem 2.1 we obtain

$$\frac{\partial V(\boldsymbol{g},n)}{\partial g_i}\ln(1+g_i) + \frac{V(\boldsymbol{g},n)}{1+g_i} = 0.$$

Substituting yields

$$Y(\boldsymbol{g},n) = \frac{\partial Y(\boldsymbol{g},n)}{\partial g_i}\ln(1+g_i)(1+g_i).$$

Note that the difference between this condition and the FOC from Theorem 2.1 (Equation (A.1)) is the derivative term. More specifically, for the optimization problem the derivative is taken with respect to the common decision variable g. Whereas for the competitive best response, the derivative is taken with respect to the decision variable under the control of the individual player, g_i. First, we consider the derivative term for the optimal case.

$$
\begin{aligned}
Y(g,n) &= \int_0^\infty (1 - F^n(x;g))dx, \\
\frac{\partial Y(g,n)}{\partial g} &= -n\int_0^\infty \frac{\partial F}{\partial g}F^{n-1}(x;g)dx.
\end{aligned}
$$

Now we consider the derivative term for the competitive case

$$
\begin{aligned}
Y(\boldsymbol{g},n) &= \int_0^\infty (1 - F(x;g_1)F(x;g_2)\cdots F(x;g_n))dx, \\
\frac{\partial Y(\boldsymbol{g},n)}{\partial g_i} &= -\int_0^\infty \frac{\partial F(x;g_i)}{\partial g}F(x;g_1)F(x;g_2)\cdots F(x;g_{i-1})F(x;g_{i+1})\cdots F(x;g_n)dx.
\end{aligned}
$$

Since the technologies are all symmetric, we know that $g = g_1 = g_2 = \cdots = g_n$. So,

$$\frac{\partial Y(g,n)}{\partial g_i} = -\int_0^\infty \frac{\partial F}{\partial g} F^{n-1}(x,g)dx.$$

Thus,

$$\frac{\partial Y(g,n)}{\partial g_i} = \frac{1}{n}\frac{\partial Y(g,n)}{\partial g}.$$

Consequently, the best response function is now

$$Y(g,n) = \frac{1}{n}Y'(g,n)\ln(1+g)(1+g).$$

If we substitute for $Y(g,n)$ based on Lemma A.1 we obtain

$$\frac{1}{n}W'(g)\ln(1+g)(1+g) - W(g) = \frac{I}{A(n)}.$$

Clearly, the presence of the $1/n$ term will shift the left hand side downward from the optimal case. Since the left hand side is strictly increasing in g (see proof of Theorem 2.1), the best response, g, for the competitive case must be greater than the optimal solution, g^*. Since the competitive policy deviates from the optimal solution, the long-run effective annual capability growth rate must be lower for the competitive case. $\qquad\square$

Proof of Theorem 2.3. We know from Theorem 2.1 that the optimal policy decreases with increasing n. If we consider the best response function from the competitive case

$$\frac{1}{n}W'(g)\ln(1+g)(1+g) - W(g) = \frac{I}{A(n)}$$

we see that both sides decrease with increasing n. Rearranging terms we obtain

$$n\left(W(g) + \frac{I}{A(n)}\right) = W'(g)\ln(1+g)(1+g).$$

It is trivial to show that the right hand side is strictly increasing. Since $A(n) \le n$ from Lemma A.1, the left hand side increases with n. Therefore, the competitive policy must increase as well, and the gap between the optimal and competitive policies must increase with n. $\qquad\square$

APPENDIX B

BASELINE INPUT PARAMETERS FOR CHAPTER 3

This appendix contains the simulation parameter values used for the base case simulation presented in Chapter 3. The parameter values are contained in the following tables.

Table B.1: General Simulation Parameters

System Types	3
Application Area Types	6
R&D Budget ($ million per year)	3000
Intersystem Delay (years)	0
Exogenous Technology Growth Rate	0.01
Internal Learning Factor	1.5

Table B.2: Technology Development Stage Parameters

Stage	Stage Costs ($ million/year)	Stage Budgets ($ million/year)	Success Probabilities (%)	Stage Length (years)
1	1	100	50	1
2	2	100	50	1
3	10	200	50	1
4	20	200	60	1
5	200	1000	70	1
6	400	1400	80	1

Table B.3: Application Area Requirements by System

	Application Areas					
Systems	1	2	3	4	5	6
1	**X**	**X**	**X**			
2				**X**	**X**	**X**
3		**X**	**X**		**X**	

Table B.4: Acquisition Life-Cycle Phase Cost Parameters($ million/year)

System	Concept Development	System Development	Production
1	20	1000	4000
2	20	1000	4000
3	20	1000	4000

Table B.5: Triangularly Distributed Parameters

Parameter	Min	Mode	Max
Base Cost Multiplier	0.5	1	2
Performance Gain Multiplier	0.8	1	1.2
Concept Development Duration (years)	2	4.9	7.5
System Development Duration (years)	1.5	2.125	8
Production Duration (years)	1.5	2	4.7

APPENDIX C

SENSITIVITY ANALYSIS FOR CHAPTER 3

A first order sensitivity analysis was conducted on the baseline parameters of the simulation presented in Chapter 3. The following figures depict the results of that analysis. Each parameter was increased and decreased by 10% from its baseline value, so the extent of each line indicates the output variation in response to that parameter. The vertical dashed lines indicate a plus 10% and minus 10% variation from the baseline output value to provide a sense of scale.

The simulation is relatively insensitive to any one parameter. The most sensitive relationship is the system operating cost to production cost. Of course, production is the largest cost in the simulation, and so it has the greatest impact. One other observation worthy of note is that, in general, capability growth rate is the most sensitive output.
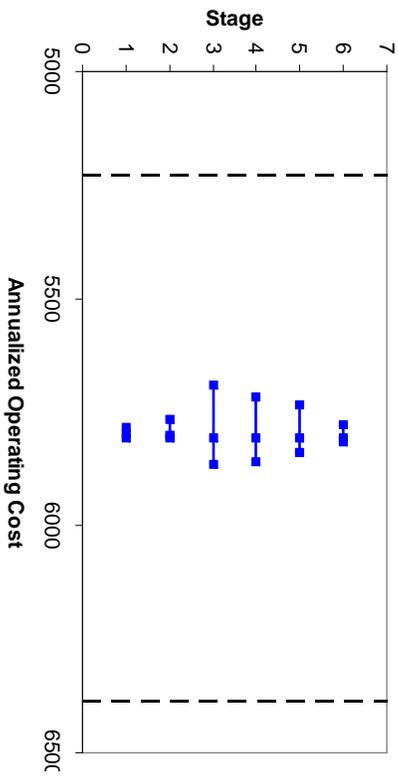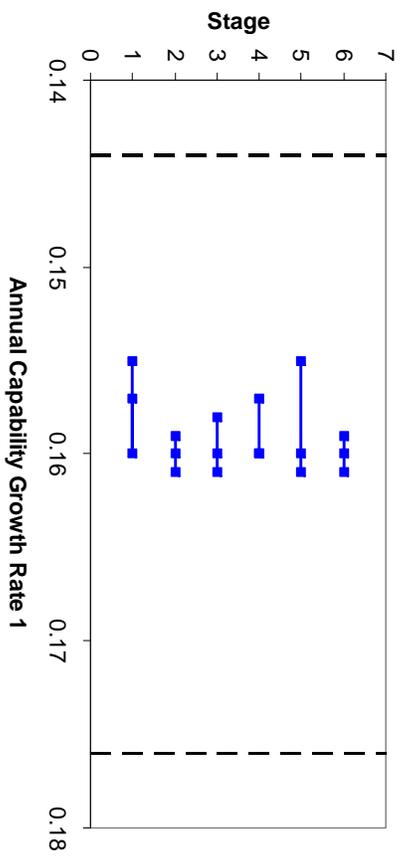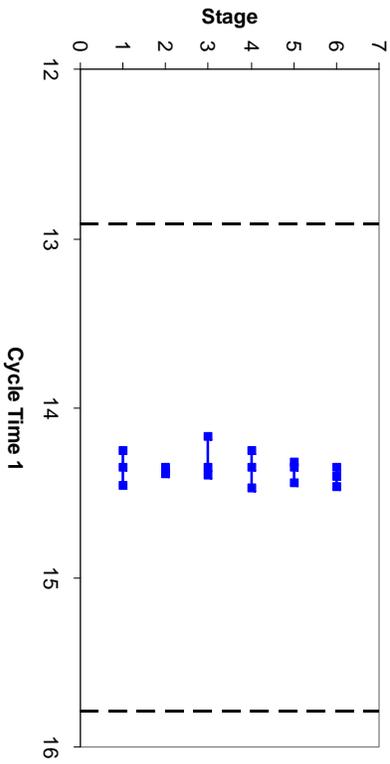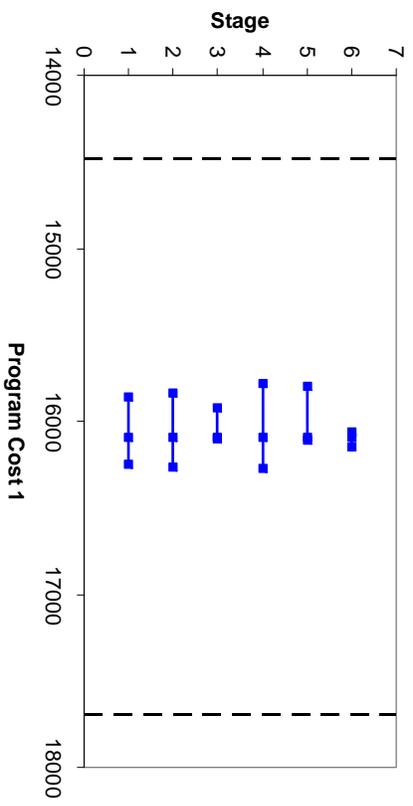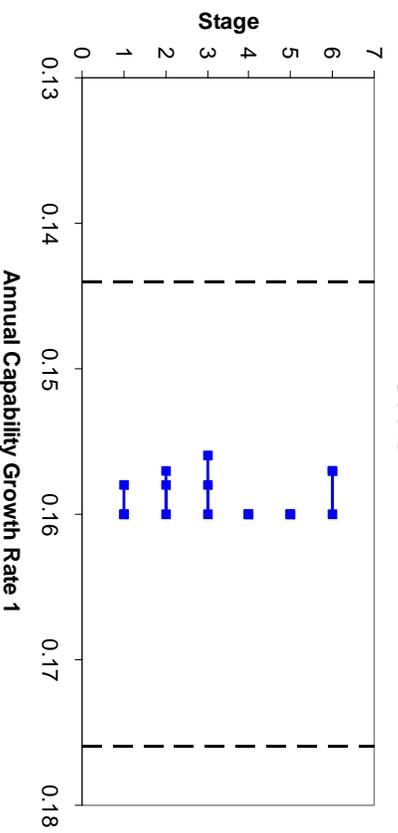
Figure C.1: Sensitivity of simulation outputs to 10% shifts in the probability of success for each technology development stage
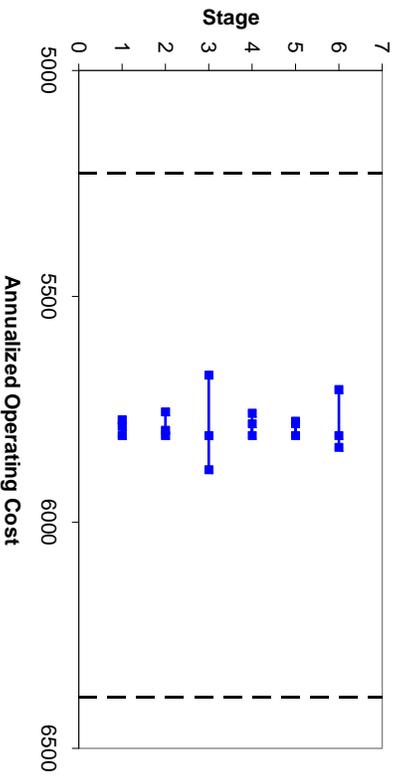
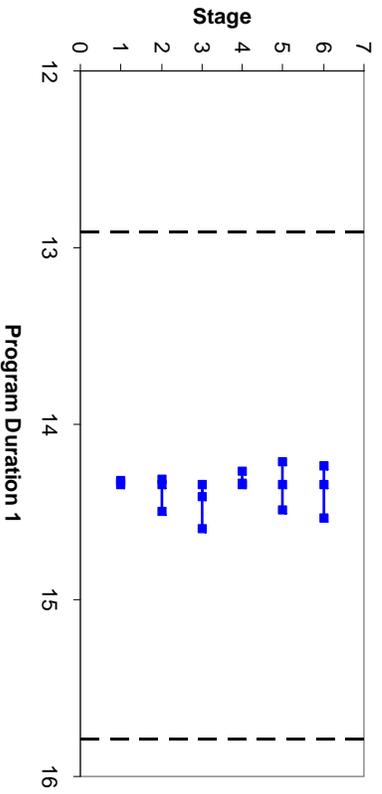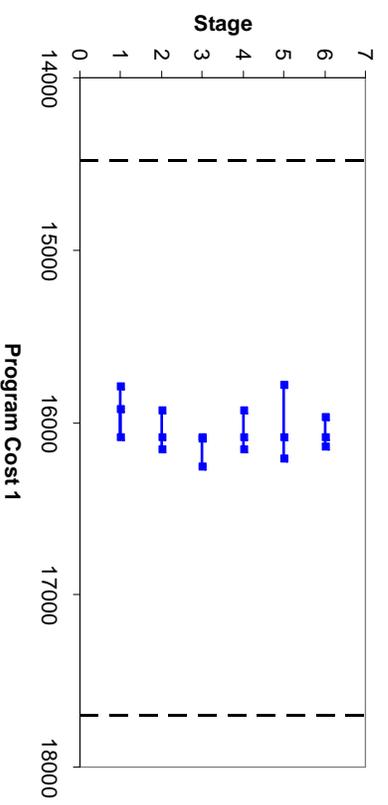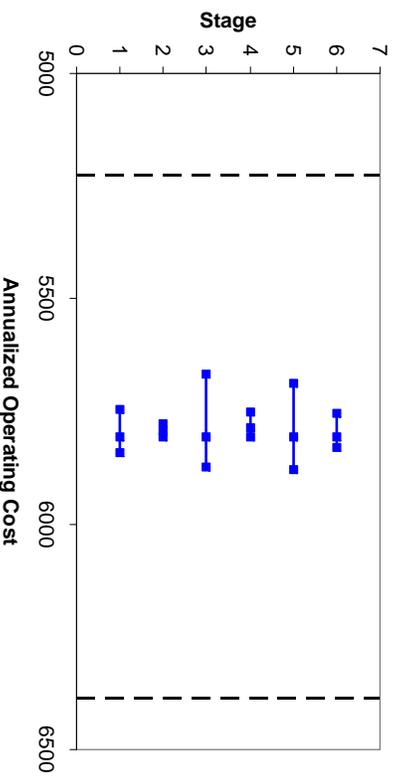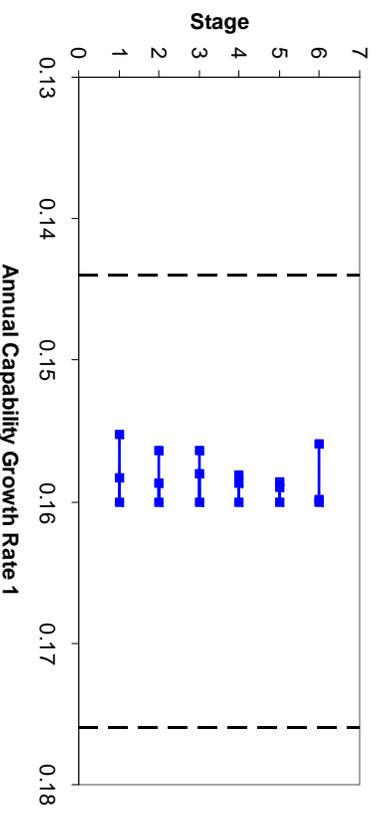**Figure C.2:** Sensitivity of simulation outputs to 10% shifts in the cost for each technology development stage

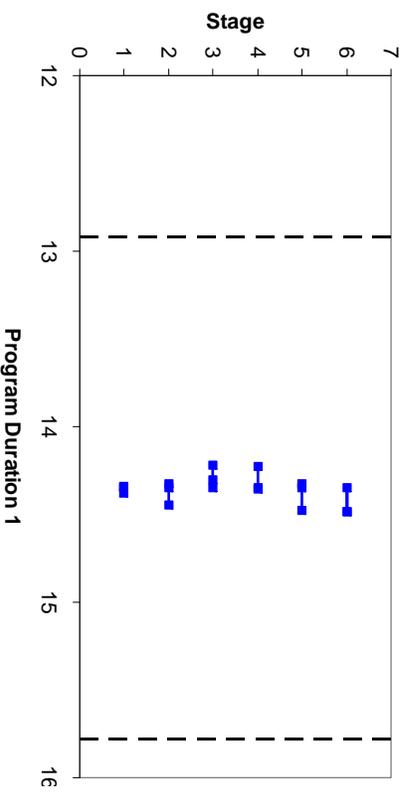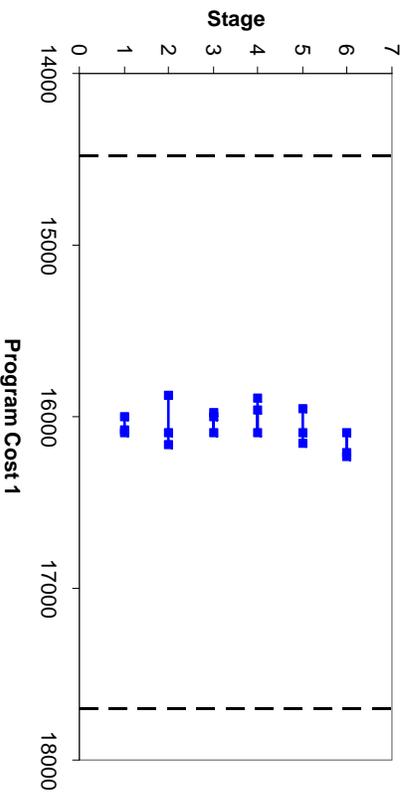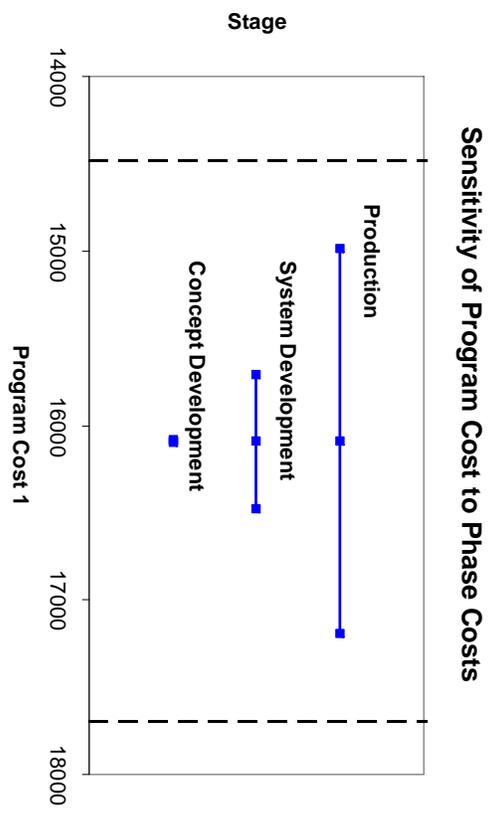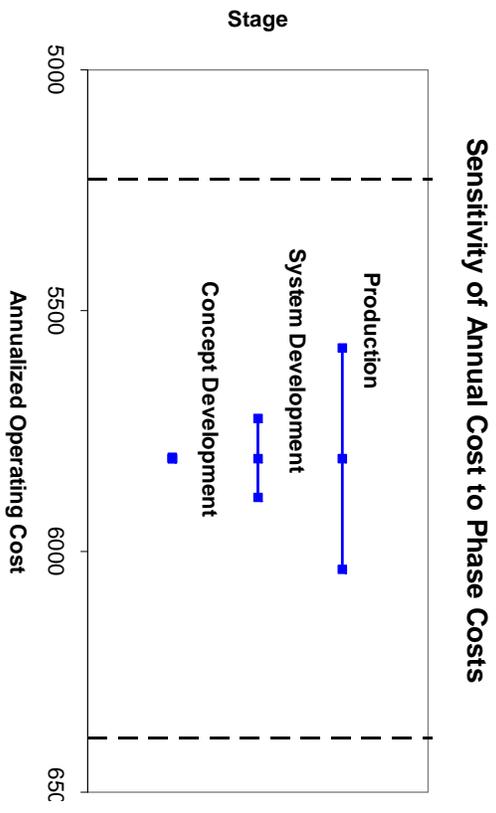**Figure C.3:** Sensitivity of simulation outputs to 10% shifts in the budget for each technology development stage

Figure C.4: Sensitivity of simulation outputs to 10% shifts in the cost for each acquisition life cycle phase

APPENDIX D

SUPPLEMENT TO CHAPTER 4

D.1   The Application of Ito's Lemma

In stochastic calculus, any transformation on an Ito process requires the application of Ito's Lemma. Since Equation (4.1) is a function of two stochastic processes, we will apply the two-dimensional version of Ito's Lemma.

$$
\begin{aligned}
dY &= \frac{\partial Y}{\partial t}dt + \frac{\partial Y}{\partial B}dB + \frac{\partial Y}{\partial C}dC + \\
&\quad \frac{1}{2}\left(\frac{\partial^2 Y}{\partial B^2}dB^2 + 2\frac{\partial^2 Y}{\partial B \partial C}dBdC + \frac{\partial^2 Y}{\partial C^2}dC^2\right)
\end{aligned}
\tag{D.1}
$$

One may note that Ito's Lemma is just the chain rule from ordinary calculus with the addition of the second derivative term. This term is required to account for the fact that Brownian motion has nonzero quadratic variation. If we substitute for dB and dC, we obtain

$$
\begin{aligned}
dY &= \frac{\partial Y}{\partial t}dt + \frac{\partial Y}{\partial B}(gBdt + \sigma_B BdZ_B) + \frac{\partial Y}{\partial C}(\alpha_C Cdt + \sigma_C CdZ_C) + \\
&\quad \frac{1}{2}\left(\frac{\partial^2 Y}{\partial B^2}\sigma_B^2 B^2 + 2\frac{\partial^2 Y}{\partial B \partial C}\rho_{BC}\sigma_B\sigma_C BC + \frac{\partial^2 Y}{\partial C^2}\sigma_C^2 C^2\right)dt
\end{aligned}
\tag{D.2}
$$

For the case of Equation (4.1) the derivative terms are

$$
\begin{aligned}
\frac{\partial Y}{\partial t} &= 0, \\
\frac{\partial Y}{\partial B} &= AaB^{a-1}C^{-a}, \\
\frac{\partial Y}{\partial C} &= -AaB^a C^{-(a+1)}, \\
\frac{\partial^2 Y}{\partial B^2} &= Aa(a-1)B^{a-2}C^{-a}, \\
\frac{\partial^2 Y}{\partial B \partial C} &= -Aa^2 B^{a-1}C^{-(a+1)}, \\
\frac{\partial^2 Y}{\partial C^2} &= Aa(a+1)B^a C^{-(a+2)}.
\end{aligned}
$$

Substituting these into Equation (D.2) and simplifying yields

$$
\begin{aligned}
dY &= \left(ag - a\alpha_C + \frac{a(a-1)}{2}\sigma_B^2 - a^2\rho_{BC}\sigma_B\sigma_C + \frac{a(a+1)}{2}\sigma_C^2 \right) Y\, dt \\
&\quad + a\sigma_B Y\, dZ_B - a\sigma_C Y\, dZ_C.
\end{aligned}
$$

D.2 Proof that $Y(t)$ is a geometric Brownian motion process

Here we will show that $Y(t)$ is a geometric Brownian motion process. First, we define the following:

$$
\begin{aligned}
\alpha_Y &= ag - a\alpha_C + \frac{a(a-1)}{2}\sigma_B^2 - a^2\rho_{BC}\sigma_B\sigma_C + \frac{a(a+1)}{2}\sigma_C^2, \\
\sigma_Y &= a\sqrt{\sigma_B^2 - 2\rho_{BC}\sigma_B\sigma_C + \sigma_C^2}, \\
dZ_Y &= \frac{a\sigma_B dZ_B - a\sigma_C dZ_C}{\sigma_Y}.
\end{aligned}
$$

Next, we need to show that dZ_Y is an increment of Brownian motion, and we can do this by finding the quadratic variation of Z_Y:

$$
\begin{aligned}
dZ_Y dZ_Y &= \frac{a^2}{\sigma_Y^2}(\sigma_B^2 dZ_B^2 - 2\sigma_B\sigma_C dZ_B dZ_C + \sigma_C^2 dZ_C^2), \\
dZ_Y dZ_Y &= \frac{a^2}{\sigma_Y^2}(\sigma_B^2 - 2\rho_{BC}\sigma_B\sigma_C + \sigma_C^2)dt, \\
dZ_Y dZ_Y &= \frac{\sigma_Y^2}{\sigma_Y^2}dt, \\
dZ_Y dZ_Y &= dt.
\end{aligned}
$$

This result implies that the quadratic variation is $[Z_Y, Z_Y](t) = t$. Thus, by the one-dimensional Lévy Theorem (see Shreve, Chapter 4 [70]), Z_Y is a Brownian motion process. Substituting terms into Equation (4.2), we obtain

$$
dY = \alpha_Y Y\, dt + \sigma_Y Y\, dZ_Y,
$$

and we see that $Y(t)$ is governed by geometric Brownian motion.

APPENDIX E

THE TRL SCALE

The Technology Readiness Level (TRL) scale is a qualitative measurement of the maturity of a technology. It has nine levels with 1 being the least mature and 9 being the most mature. The scale was originally developed by NASA [48], but has since been adopted by the Department of Defense. The following table presents the defense version of the scale.

Table E.1: DoD Technology Readiness Levels (Source: Defense Acquisition Guidebook, Chapter 10.5.2 [20])

Technology Readiness Level	Description
1. Basic principles observed and reported.	Lowest level of technology readiness. Scientific research begins to be translated into applied research and development. Examples might include paper studies of a technology's basic properties.
2. Technology concept and/or application formulated.	Invention begins. Once basic principles are observed, practical applications can be invented. Applications are speculative and there may be no proof or detailed analysis to support the assumptions. Examples are limited to analytic studies.
3. Analytical and experimental critical function and/or characteristic proof of concept.	Active research and development is initiated. This includes analytical studies and laboratory studies to physically validate analytical predictions of separate elements of the technology. Examples include components that are not yet integrated or representative.
4. Component and/or breadboard validation in laboratory environment.	Basic technological components are integrated to establish that they will work together. This is relatively "low fidelity" compared to the eventual system. Examples include integration of "ad hoc" hardware in the laboratory.
5. Component and/or breadboard validation in relevant environment.	Fidelity of breadboard technology increases significantly. The basic technological components are integrated with reasonably realistic supporting elements so it can be tested in a simulated environment. Examples include "high fidelity" laboratory integration of components.
6. System/subsystem model or prototype demonstration in a relevant environment.	Representative model or prototype system, which is well beyond that of TRL 5, is tested in a relevant environment. Represents a major step up in a technology's demonstrated readiness. Examples include testing a prototype in a high-fidelity laboratory environment or in simulated operational environment.
7. System prototype demonstration in an operational environment.	Prototype near, or at, planned operational system. Represents a major step up from TRL 6, requiring demonstration of an actual system prototype in an operational environment such as an aircraft, vehicle, or space. Examples include testing the prototype in a test bed aircraft.
8. Actual system completed and qualified through test and demonstration.	Technology has been proven to work in its final form and under expected conditions. In almost all cases, this TRL represents the end of true system development. Examples include developmental test and evaluation of the system in its intended weapon system to determine if it meets design specifications.
9. Actual system proven through successful mission operations.	Actual application of the technology in its final form and under mission conditions, such as those encountered in operational test and evaluation. Examples include using the system under operational mission conditions.

REFERENCES

[1] ARENA, M. V., BLICKSTEIN, I., YOUNOSSI, O., and GRAMMICH, C. A., *Why has the Cost of Navy Ships Risen?: A Macroscopic Examination of the Trends in U.S. Naval Ship Costs Over the Past Several Decades.* Santa Monica, CA: RAND, 2006.

[2] ARENA, M. V., LEONARD, R. S., MURRAY, S. E., and YOUNOSSI, O., "Historical cost growth of completed weapon system programs," Tech. Rep. TR-343-AF, RAND, Santa Monica, CA, 2006.

[3] BALDWIN, W. L., *The Structure of the Defense Market, 1955–1964.* Durham, NC: Duke University Press, 1967.

[4] BIRKLER, J., RUSHWORTH, D., CHIESA, J., PUNG, H., ARENA, M. V., and CHANK, J. F., *Differences Between Military and Commercial Shipbuilding.* Santa Monica, CA: RAND, 2005.

[5] BLACK, F. and SCHOLES, M., "The pricing of options and corporate liabilities," *Journal of Political Economy*, vol. 81, pp. 637–654, 1973.

[6] BODNER, D. A. and ROUSE, W. B., "Understanding R&D value creation with organizational simulation," *Systems Engineering*, vol. 10, no. 1, pp. 64–82, 2007.

[7] BOUDREAU, M. W., "Acousitc rapid COTS insertion-case study," in *Proceedings of the Third Annual Acquisition Research Symposium*, (Monterey), pp. 185–186, Naval Postgraduate School, 2006.

[8] BURGESS, G. M., *A Management Analysis and Systems Model of Department of Defense Acquistion Structure and Policy.* PhD thesis, The Florida State University, 1988.

[9] CANCIAN, M., "Acquisition reform: It's not as easy as it seems," *Acquisition Review Quarterly*, pp. 190–192, Summer 1995.

[10] CHRISTENSEN, D. S., SEARLE, D. A., and VICKERY, C., "The impact of the Packard Commission's recommendations on reducing cost overruns on defense acquisition contracts," *Acquisition Review Quarterly*, pp. 252–256, Summer 1999.

[11] CLARK, JR., T. J., WHITTENBERG, E. L., and WOODRUFF, A. H., "A policy analysis of defense weapons acquisition," *Policy Studies Journal*, vol. 14, no. 1, pp. 21–46, 1985.

[12] CLARKSON, P. J., SIMONS, C., and ECKERT, C., "Predicting change propagation in complex design," *Transactions of the ASME*, vol. 126, pp. 788–797, 2004.

[13] COOK, C. R. and GRASER, J. C., *Military Airframe Acquisition Costs: The Effects of Lean Manufacturing.* Santa Monica, CA: RAND, 2001.

[14] CORNFORD, S. L. and SARSFIELD, L., "Quantitative methods for maturing and infusing advanced spacecraft technology," in *IEEE Aerospace Conference Proceedings*, pp. 663–681, IEEE, 2004.

[15] COX, J. C., ROSS, S. A., and RUBINSTEIN, M., "Option pricing: A simplified approach," *Journal of Financial Economics*, vol. 7, pp. 229–263, 1979.

[16] DILLARD, J. and FORD, D., "Too little too soon? Modeling the risks of spiral development," in *Proceedings of the Fourth Annual Acquisition Research Symposium*, (Monterey, CA), pp. 484–500, Naval Postgraduate School, May 2007.

[17] DIXIT, A. K. and PINDYCK, R. S., *Investment Under Uncertainty.* Princeton: Princeton University Press, 1994.

[18] DoD, "DoD Directive 5000.1," May 12 2003.

[19] DoD, "DoD Instruction 5000.2," May 12 2003.

[20] DoD, "Defense acquisition guidebook." http://akss.dau.mil/dag/welcome.asp, 2006. Retrieved November 12, 2007.

[21] DoD, "Seleced acquisition report(SAR) summary tables," April 6 2007.

[22] DoD, "Seleced acquisition report(SAR) summary tables," August 14 2007.

[23] DREZNER, J. A., JARVAISE, J., HESS, R., HOUGH, P., and NORTON, D., *An Analysis of Weapon System Cost Growth (MR-291-AF).* Santa Monica, CA: RAND, 1993.

[24] ECKERT, C., CLARKSON, P. J., and ZANKER, W., "Change and customization in complex engineering domains," *Research in Engineering Design*, vol. 15, pp. 1–21, 2004.

[25] FOX, J. R., *Arming America: How the U.S. Buys Weapons.* Boston: Harvard University Press, 1974.

[26] FOX, J. R., *The Defense Management Challenge: Weapons Acquisition.* Boston: Havard Business School Press, 1988.

[27] GANSLER, J. S. and LUCYSHYN, W., "A strategy for defense acquisition research." Center for Public Policy and Private Enterprise, University of Maryland, August 2005.

[28] GANSLER, J. S., *The Defense Industry.* Cambridge: The MIT Press, 1980.

[29] GAO, "Acquisitions: DoD's defense acquisition improvement program: A status report." United States Government Accountability Office, GAO/NSIAD-86-148, July 1986.

[30] GAO, "Defense acquisitions: DoD's revisied policy emphasizes best practices, but more controls are needed." United States Government Accountability Office, GAO-04-53, November 10, 2003.

[31] GAO, "Best practices: Stronger practices needed to improve DOD technology transition processes." United States Government Accountability Office, GAO-06-883, September 5, 2006.

[32] GAO, "Defense acquisitions: Actions needed to get better results on weapon systems investments." United States Government Accountability Office, GAO-06-585T, April 5, 2006.

[33] GAO, "Defense acquisitions: Major weapon systems continue to experience cost and schedule problems under DoD's revised policy." United States Government Accountability Office, GAO-06-368, April 13, 2006.

[34] GAO, "Defense contracting-questions for the record." United States Government Accountability Office, GAO-07-217R, December 21, 2006.

[35] GAO, "Tactical aircraft: Questions concerning the F-22A's business case." United States Government Accountability Office, GAO-06-991T, July 25, 2006.

[36] GAO, "Best practices: An integrated portfolio management approach to weapon system investments could improve DOD's acquisition outcomes." United States Government Accountability Office, GAO-07-388, March 30, 2007.

[37] GAO, "Defense acquisitions: Assessments of selected weapons programs." United States Government Accountability Office, GAO-07-406SP, March 30 2007.

[38] HANKS, C. H., AXELBAND, E. I., LINDSAY, S., MALIK, M. R., and STEELE, B. D., *Reexamining Military Acquisition Reform: Are We There Yet?* Santa Monica, CA: RAND, 2005.

[39] HAYES, R. H. and ABERNATHY, W. J., "Managing our way to economic decline," *Harvard Business Review*, vol. 58, no. 4, pp. 67–77, 1980.

[40] HAYES, R. H. and GARVIN, D. A., "Managing as if tomorrow mattered," *Harvard Business Review*, vol. 60, no. 3, pp. 70–79, 1982.

[41] HOLLAND, L., *Weapons Under Fire.* New York: Garland Publishing, 1997.

[42] HOUSEL, T., MUN, J., and TARANTINO, E., "AEGIS and ship self-defense system (SSDS) platforms: Using KVA analysis, risk simulation and strategic real options to assess operational effectiveness," in *Proceedings of the Fourth Annual Acquisition Research Symposium*, (Monterey, CA), pp. 122–162, Naval Postgraduate School, May 2007.

[43] JOHNSON, W. M. and JOHNSON, C. O., "The promise and perils of spiral acquistion: A practical approach to evolutionary acquisition," *Acquisition Review Quarterly*, pp. 175–188, Summer 2002.

[44] KAMRAD, B. and RITCHKEN, P., "Multinomial approximating models for options with k state variables," *Management Science*, vol. 37, no. 12, pp. 1640–1652, 1991.

[45] KIRN, M., "Examining NASA's use of technology readiness levels," tech. rep., NASA, August 2005.

[46] LEONARD, R. S., DREZNER, J. A., and SOMMER, G., *The Arsenal Ship Acqusition Process Experience*. Santa Monica, CA: RAND, 1999.

[47] LORELL, M. A., LOWELL, J. F., and YOUNOSSI, O., *Evolutionary Acquisition: Implementation Challenges for Defense Space Programs*. Santa Monica, CA: RAND, 2006.

[48] MANKINS, J. C., "Technology readiness levels: A white paper," April 6, 1995. http://www.hq.nasa.gov/office/codeq/trl/.

[49] MARKOWITZ, H. M., "Portfolio selection," *Journal of Finance*, vol. 7, no. 1, pp. 77–91, 1952.

[50] MARKOWITZ, H. M., "The optimization of a quadratic function subject to linear constraints," *Naval Research Logistics Quarterly*, vol. 3, no. 1-2, pp. 111–133, 1956.

[51] MARKOWITZ, H. M., *Mean-Variance Analysis in Portfolio Choice and Capital Markets*. New York: Basil Blackwell, 1987.

[52] MARKOWITZ, H. M., *Portfolio Selection*. New York: Wiley, 1987.

[53] MEYERS, S. C., "Determinants of corporate borrowing," *Journal of Financial Economics*, vol. 5, pp. 147–175, 1977.

[54] MODIGLIANI, F. and MILLER, M. H., "The cost of capital, corporate finance, and the theory of investment," *American Economic Review*, vol. 48, no. 3, pp. 261–297, 1958.

[55] MOORHOUSE, D. J., "Detailed definitions and guidance for application of technology readiness levels," *Journal of Aircraft*, vol. 39, no. 1, pp. 190–192, 2002.

[56] PECK, M. J. and SCHERER, F. M., *The Weapons Acquisition Process: An Economic Analysis*. Cambridge, MA: Harvard University Press, 1962.

[57] PENNOCK, M. J., ROUSE, W. B., and KOLLAR, D. L., "Transforming the acquisition enterprise: A framework for analysis and a case study of ship acquisition," *Systems Engineering*, vol. 10, no. 2, pp. 99–117, 2007.

[58] POLLAK, R. A., "The theory of the cost of living," in *Price Level Measurement* (DIEWERT, W. E., ed.), Amsterdam: Elsevier, 1990.

[59] PRESIDENT'S BLUE RIBBON COMMISSION ON DEFENSE, "A quest for excellence," 1986. accessed: http://www.ndu.edu/library/pbrc/pbrc.html.

[60] REED, T. S., COWDREY, J., and PIKE, W., "Monsters in the closet? The impact of collective bargaining agreements in A-76 sourcing decisions," in *Proceedings of the Fourth Annual Acquisition Research Symposium*, (Monterey, CA), pp. 314–329, Naval Postgraduate School, May 2007.

[61] ROGERS, E. W. and BIRMINGHAM, R. P., "A ten-year review of the vision for transforming the defense acquisition system," *Defense Acquisition Review Journal*, pp. 37–61, January-April 2004.

[62] ROUSE, W. B., "Enterprises as systems: Essential challenges and approaches to transformation," *Systems Engineering*, vol. 8, no. 2, pp. 138–150, 2005.

[63] ROUSE, W. B., "A theory of enterprise transformation," *Systems Engineering*, vol. 8, no. 4, pp. 279–295, 2005.

[64] ROUSE, W. B., ed., *Enterprise Transformation: Understanding and Enabling Fundamental Change*. New York: Wiley, 2006.

[65] ROUSE, W. B., PENNOCK, M. J., and KOLLAR, D. L., "Transforming the enterprise of acquiring public sector complex systems," in *Proceedings of the Third Annual Acquisition Research Symposium*, (Monterey, CA), pp. 113–127, 2006.

[66] SAMMET, JR., G. and GREEN, D. E., *Defense Acquisition Management*. Boca Raton: Florida Atlantic University Press, 1990.

[67] SANDLER, T. and HARTLEY, K., *The Economics of Defense*. Cambridge: Cambridge University Press, 1995.

[68] SCHERER, F. M., *The Weapons Acquisition Process: Economic Incentives*. Cambridge, MA: Harvard University Press, 1964.

[69] SHAPIRO, A. A., "Technology infusion for space-flight programs," in *IEEE Areospace Conference Proceedings*, pp. 653–662, IEEE, 2004.

[70] SHREVE, S. E., *Stochastic Calculus for Finance II: Continuous-Time Models*. New York: Springer, 2004.

[71] SMALING, R. and DE WECK, O., "Assessing risks and opportunities of technology infusion in system design," *Systems Engineering*, vol. 10, no. 1, pp. 1–25, 2007.

[72] STEVENS, G. A. and BURLEY, J., "3,000 raw ideas = 1 commerical success!," *Research Technology Management*, vol. 40, no. 3, pp. 16–27, 1997.

[73] VARIAN, H. R., *Microeconomic Analysis.* New York: W. W. Norton & Company, 3rd ed., 1992.

[74] WEIDA, W. J. and GERTCHER, F. L., *The Political Economy of National Defense.* Boulder, CO: Westview Press, 1987.