



EXCERPT FROM THE PROCEEDINGS

OF THE
FOURTH ANNUAL ACQUISITION
RESEARCH SYMPOSIUM
WEDNESDAY SESSIONS

**Development vs. Deployment: How Mature Should a Technology be
before it is Considered for Inclusion in an Acquisition Program?**

Published: 30 April 2007

by

Michael J. Pennock, Research Fellow,

William B. Rouse, PhD, Executive Director and Professor, and

Diane Kollar, Director, Industry and Government Relations,

Tennenbaum Institute

**4th Annual Acquisition Research Symposium
of the Naval Postgraduate School:**

**Acquisition Research:
Creating Synergy for Informed Change**

May 16-17, 2007

Approved for public release, distribution unlimited.

Prepared for: Naval Postgraduate School, Monterey, California 93943



The research presented at the symposium was supported by the Acquisition Chair of the Graduate School of Business & Public Policy at the Naval Postgraduate School.

To request Defense Acquisition Research or to become a research sponsor, please contact:

NPS Acquisition Research Program
Attn: James B. Greene, RADM, USN, (Ret)
Acquisition Chair
Graduate School of Business and Public Policy
Naval Postgraduate School
555 Dyer Road, Room 332
Monterey, CA 93943-5103
Tel: (831) 656-2092
Fax: (831) 656-2253
E-mail: jbgreene@nps.edu

Copies of the Acquisition Sponsored Research Reports may be printed from our website www.acquisitionresearch.org

Conference Website:
www.researchsymposium.org



Acquisition Research program
GRADUATE SCHOOL OF BUSINESS & PUBLIC POLICY
NAVAL POSTGRADUATE SCHOOL

Proceedings of the Annual Acquisition Research Program

The following article is taken as an excerpt from the proceedings of the annual Acquisition Research Program. This annual event showcases the research projects funded through the Acquisition Research Program at the Graduate School of Business and Public Policy at the Naval Postgraduate School. Featuring keynote speakers, plenary panels, multiple panel sessions, a student research poster show and social events, the Annual Acquisition Research Symposium offers a candid environment where high-ranking Department of Defense (DoD) officials, industry officials, accomplished faculty and military students are encouraged to collaborate on finding applicable solutions to the challenges facing acquisition policies and processes within the DoD today. By jointly and publicly questioning the norms of industry and academia, the resulting research benefits from myriad perspectives and collaborations which can identify better solutions and practices in acquisition, contract, financial, logistics and program management.

For further information regarding the Acquisition Research Program, electronic copies of additional research, or to learn more about becoming a sponsor, please visit our program website at:

www.acquisitionresearch.org

For further information on or to register for the next Acquisition Research Symposium during the third week of May, please visit our conference website at:

www.researchsymposium.org



THIS PAGE INTENTIONALLY LEFT BLANK



Development vs. Deployment: How Mature Should a Technology be Before it is Considered for Inclusion in an Acquisition Program?

Presenter: Michael Pennock is a research fellow for the Tennenbaum Institute for Enterprise Transformation as well as a PhD candidate in Industrial and Systems Engineering at Georgia Tech. He has previously worked as a systems engineer for the Northrop Grumman Corporation, and he earned his Bachelor's and Master's degrees in Systems Engineering from the University of Virginia. His research focuses on adapting economic analysis to address problems in systems engineering.

Author: Bill Rouse is the Executive Director of the Tennenbaum Institute at the Georgia Institute of Technology. He is also a professor in the College of Computing and School of Industrial and Systems Engineering. Rouse has written hundreds of articles and book chapters, and has authored many books, including most recently *People and Organizations: Explorations of Human-Centered Design* (Wiley, 2007), *Essential Challenges of Strategic Management* (Wiley, 2001) and the award-winning *Don't Jump to Solutions* (Jossey-Bass, 1998). He is editor of *Enterprise Transformation: Understanding and Enabling Fundamental Change* (Wiley, 2006), co-editor of *Organizational Simulation: From Modeling & Simulation to Games & Entertainment* (Wiley, 2005), co-editor of the best-selling *Handbook of Systems Engineering and Management* (Wiley, 1999), and editor of the eight-volume series *Human/Technology Interaction in Complex Systems* (Elsevier). Among many advisory roles, he has served as Chair of the Committee on Human Factors of the National Research Council, a member of the US Air Force Scientific Advisory Board, and a member of the DoD Senior Advisory Group on Modeling and Simulation. Rouse is a member of the National Academy of Engineering, as well as a fellow of four professional societies— Institute of Electrical and Electronics Engineers (IEEE), the International Council on Systems Engineering (INCOSE), the Institute for Operations Research and Management Science, and the Human Factors and Ergonomics Society.

Author: Diane Kollar is Director of Industry and Government Relations for the Tennenbaum Institute at the Georgia Institute of Technology. Prior to this position, she was the Director of Development for the School of Industrial and Systems Engineering at Georgia Tech. She has held several positions at Georgia Tech before which she was Associate Director of Development at The Carter Center. She has served in various other development roles in a range of nonprofits, including positions in corporate relations and public relations. Her interests and expertise include resource-development strategy formulation and organizational implementation, particularly in public sector and nonprofit enterprises, as well as public policy issues associated with such strategies and organizations. Ms. Kollar received her BA in Government and International Studies and Master of Public Administration from the University of South Carolina. She also attended the Bryce Harlow Institute on Business and Government Affairs at Georgetown University and studied organizational behavior at Florida Atlantic University.

Michael J. Pennock
Research Fellow
Tennenbaum Institute
755 Ferst Drive
Atlanta, GA 30332-0205
mpennock@ti.gatech.edu

William B. Rouse, PhD
Executive Director and Professor
Tennenbaum Institute
755 Ferst Drive
Atlanta, GA 30332-0205



(404) 894-2331
bill.rouse@ti.gatech.edu

Diane Kollar
Director, Industry and Government Relations
Tennenbaum Institute
755 Ferst Drive
Atlanta, GA 30332-0205
(404) 894-7014
diane.kollar@ti.gatech.edu

Abstract

Modern military systems increasingly rely on the integration of multiple advanced technologies. While these technologies vastly increase warfighter capabilities, they also introduce risk into the system design and development process that tends to increase both its cost and duration. As acquisition cycle-times increase, warfighters must make do with dated technology for longer periods. Thus, there is an incentive to push as many advanced technologies as possible into each program to maximize warfighter capability over the next acquisition cycle. Unfortunately, the more new technologies a system has, the more risky its acquisition becomes, and consequently, its duration and cost increase even further. Thus, there is a feedback effect that exacerbates the problem. Open-architecture designs can partially alleviate this problem, but some technology decisions are so integral to a system's design that they cannot be relegated to future upgrades. Consequently, there is a tradeoff between incorporating these technologies now and increasing program risk or developing and evaluating them further but potentially postponing their application to future acquisition cycles. Our paper will examine this tradeoff by considering a new technology's contribution to program risk.

Introduction

Despite repeated attempts at reforming the defense acquisition process, Defense Department programs continue to experience substantial cost overruns, schedule delays, and performance shortfalls. While there are likely multiple causes for reform failure, this paper aims to address only one of the critical issues that contribute to these acquisition challenges. That issue is the maturity of critical technologies employed in major defense acquisition programs.

There have been repeated calls for the Department of Defense to use evolutionary rather than revolutionary acquisition strategies. In fact, the DoD has revised its acquisition policies to that end (GAO, 2003). Despite these new policies, recent 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 (GAO, 2006, April 5; 2006, April 13; 2006, December 21). It seems that every program is an exception. Why is this?

To that end, this paper investigates two key questions: What level of maturity is acceptable for a technology to be included in a major acquisition program, and what obstacles prevent the DoD from implementing an evolutionary acquisition process?

Our findings will show that, relatively speaking, it is better to employ mature technologies; thus, an evolutionary strategy is superior under most circumstances to a



revolutionary strategy in terms of getting capabilities delivered to the warfighter. We also found, however, that when a program relies on multiple, critical technologies, especially those intended for a multi-mission role, the evolutionary strategy is unstable. There is a natural tendency to revert to the revolutionary technology strategy even though it is not in the best interest of the warfighter.

This paper is structured in the following manner. First, we discuss the background of knowledge-based, evolutionary acquisition and why it is considered important for defense acquisition. Second, we develop a high-level simulation model of acquisition to help us investigate these issues. Third, we use the model to analyze defense acquisition policy alternatives regarding technological maturity. Finally, we conclude with the policy implications of this analysis.

Background

The troubled history of the DoD acquisition system (as well as the repeated attempts to reform it) are well known, and we will not recount them here (See Pennock, Rouse & Kollar, 2007 and GAO, 2006, April 13). Instead, our focus will be on the more recent attempts to reform the acquisition system by employing knowledge-based business practices and evolutionary acquisition.

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 in which programs must meet certain requirements before proceeding from one phase to the next (DOD, 2003a, 2003b). (See Figure 1.) Part of this milestone process is an assessment of the maturity of technologies to be employed in acquisition programs as well as a plan to manage their development.

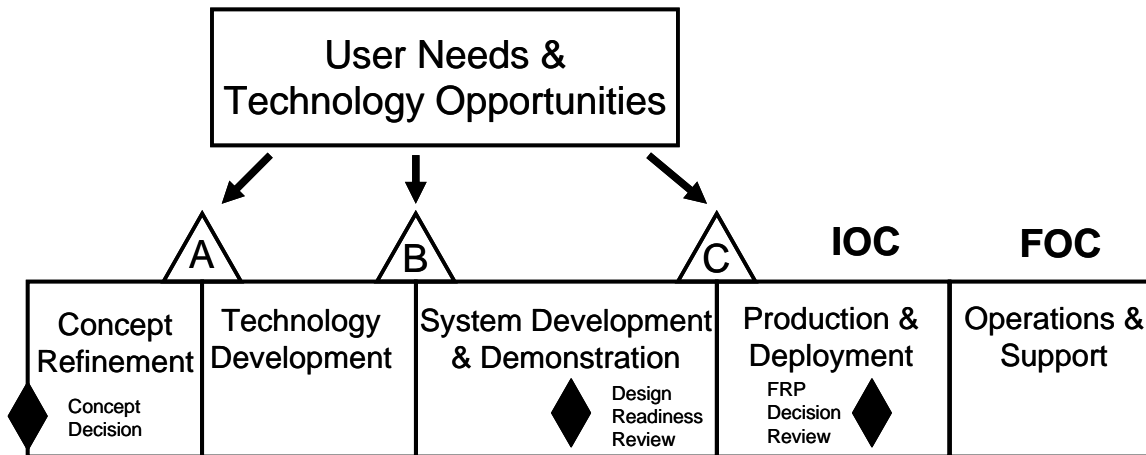


Figure 1. Defense Acquisition Management Framework
(DoD, 2003b)

Technological maturity is typically assessed using the Technology Readiness Level (TRL) scale (Table 1). 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 (Smailing & deWeck 2007). 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, qualification to enter 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 (GAO 2006, April 13)).

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.

Table 1. DoD Technology Readiness Levels
(DoD, 2006, Ch. 10.5.2)

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, these tend 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, 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 often referred 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 (Johnson & Johnson, 2002). 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 (Boudreau, 2006).

When technologies or subsystems are tightly coupled to the overall system, however, any changes to the design of the subsystem impact the design of the whole system. Thus, open-architecture design is not always a feasible alternative. An extreme example would be the hull-form of a surface combatant. Take, for instance, the tumblehome hull design of the new Zumwalt-class destroyer. 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 design process 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 policies (GAO, 2006, April 5; 2006, April 13; 2006, December 21). 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 (GAO 2006, April 13).

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 (GAO 2006, December 21). 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. (GAO 2006, December 21)

Thus, a more evolutionary approach to acquisition may have reduced the risks to the warfighter by both 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 for this, one of which is the lack of mandatory controls on the milestone process (GAO 2003, 2006, April 5, 2006, April 13, 2006, December 21). But if evolutionary acquisition is superior, why would the DoD not follow its tenets even without the mandatory controls? There are really two possibilities. Either evolutionary acquisition is not the best approach and when given the flexibility program managers avoid it, or the nature of the acquisition system itself works against the successful implementation of evolutionary methods.



Analysis Approach

To better understand the nature of evolutionary acquisition, we must model the impact of a program's technology strategy on the level of capability actually deployed in the field. In particular, a technology strategy consists of the technologies selected to improve each capability that a system provides. A technology policy that emphasizes major increases in capability would likely rely on immature technology and, thus, would be a revolutionary strategy. Consequently, the acquisition program will require a substantial technology development phase. On the other hand, a technology strategy that emphasizes small improvements in technology would rely on more mature technology and could be considered an evolutionary strategy. This type of strategy effectively detaches technology development from the acquisition program and, consequently, would have a relatively short technology-development phase.

What we would like to examine is the impact of the selected technology strategy over the long-term. Thus, we are concerned with the deployed capability resulting from a sequence of acquisition programs. In particular, we are assuming that our objective is to improve the capabilities of a particular class of system such as a surface combatant or air superiority fighter. To model this, we must establish a means to link the selected technology strategy to the time required to complete an acquisition program. This will determine when a capability improvement is deployed. After an acquisition program completes, we assume that another begins immediately to procure the next iteration of that system.

To accomplish this, we will assume that we can model each acquisition program as a small PERT chart. PERT charts are a common program management tool for managing schedule risk. For our particular model, we will assume a fairly simple formulation. We will assume that there is a technology-development stage followed by a system-integration stage. Each acquisition program contains a number of critical technologies that must be developed for the program to reach a successful conclusion. We will assume 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, since unanticipated outcomes in the technology-development phase can lead to substantial rework in the integration phase, this is not an unreasonable assumption. Given those assumptions, we can structure each acquisition program as shown in Figure 2.






Figure 2. Simplified PERT Chart for an Acquisition Program

Keeping in line with the standard PERT formulation, we will assume that the duration of each technology development activity is stochastic and beta distributed (Figure 3). The beta distribution is appealing in this context because it has finite upper and lower bounds on the activity duration, hence its use in PERT.




Figure 3. Example Beta Distribution ($\alpha = 2$, $\beta = 5$)

One notion we would like to capture is the relationship between the maturity of a technology selected for an acquisition program and the amount of schedule risk it entails. It is fairly safe to assume that the more immature a technology is, the more schedule risk there is in its development. In fact, we can go one step further and assume that it follows the law of diminishing returns. In other words, each additional increment of schedule risk that we

accept buys a reduced amount of gain in capability. To make this relationship more concrete, we must select metrics for the gain in capability and the level of schedule risk. For the former, we will consider the percent gain in capability over the currently deployed capability. Thus, a relatively low percent gain would be considered an evolutionary technology whereas a large percent gain would be a revolutionary technology. Since we would only accept an immature technology in exchange for an increase in capability, we can assume that for the purposes of our model, the percent gain in capability is also an acceptable proxy for technological maturity. As for the risk, we will assume that schedule risk is encapsulated in the upper bound of the probability distribution for the duration of technology development. For the sake of simplicity, the lower bound and shape parameters of the beta distribution will remain constant. Thus, if we select a particular percent gain in capability as our technology policy, it determines a particular upper bound on the distribution of the development time of that technology. This is illustrated in Figure 4. When we change the upper bound of the distribution, two things occur. We increase the expected time to develop the technology, and we increase the spread of the distribution.




Figure 4. Tradeoff between Risk (the Upper Bound on the Duration of Technology Development) and Return (the Growth in Capability)

We define a technology policy as the targeted percent gain in capability for each acquisition program. Thus, if a more aggressive target is selected, there will be a greater increase in capability for each new system deployed. However, the expected duration of the acquisition cycle will also increase. Given our model structure, we can use Monte Carlo simulation to generate possible capability trajectories. This is accomplished in the following manner. First, sample from the beta distribution for each technology is included in the acquisition program. The integration phase cannot begin until all technology development is complete, so the longest sampled time dictates the length of the technology-development phase. That time plus the time required for integration is the total time required for the acquisition program. At the end of the acquisition program, each capability is increased by the gain targeted in the technology policy. The process then repeats again with the next

acquisition program. This yields a capability trajectory for the technology policy. One example is shown in Figure 5.

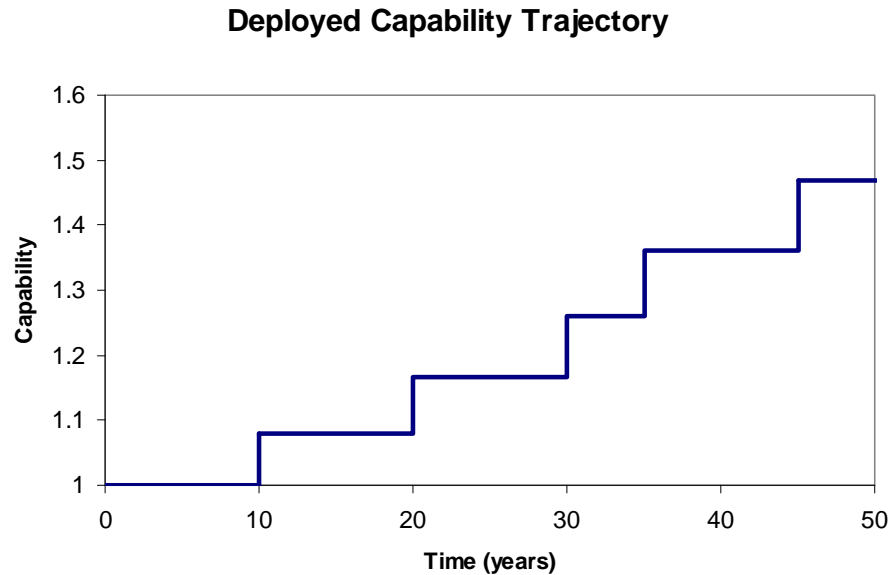


Figure 5. Sample Capability Trajectory

We see in Figure 5 that, with our model, the capability trajectory is a step function because of the discrete nature of acquisition programs. Thus, we see that the longer the acquisition cycle, the longer warfighters must make due with older equipment. To facilitate analysis, we would like to capture the value of any given capability trajectory as a single number. We will do so through the average deployed capability. To calculate the average deployed capability, we select a time horizon, say 50 years, and then calculate the average value of capability over that time interval. While this is not a perfect metric, the notion we are trying to capture is the level of capability that warfighters can expect from their equipment if they are forced to engage in hostilities without warning. This allows us to compare the competing strategies of small-but-rapid capability increments versus large-but-infrequent capability increments. If we generate many sample capability trajectories for a particular technology policy, we can calculate the expected average deployed capability to evaluate the efficacy of that policy.

Analysis Results

The first question we would like to consider is whether it is better to pursue an evolutionary vs. a revolutionary technology strategy. From a purely performance standpoint, we can answer this question using the model we described above but with only a single technology for each acquisition program. To make this more concrete, we will assume some parameter values and run our Monte Carlo simulation over a range of technology policies. In particular, the range we consider is a capability gain per acquisition cycle

between 10 and 30%. The relationship to the upper bound of the duration distribution is described in Figure 4. This is just the function:

$$\text{Function 1. Capability Gain} = 0.071841 (\text{Upper Bound})^{0.477121}$$

Of course, other functional forms are possible, and we will discuss these in more detail later. Under this function, the upper bound of the resulting beta distribution can vary between 2 and 20 years. As far as defining the rest of the beta distribution, the lower bound is always 2 years, and the shape parameters are $\alpha = 2$ and $\beta = 5$. For the purposes of calculating the average deployed capability, the initial level of capability is always one, and the time horizon is 50 years. To emphasize the impact of technology development, we will assume that the duration of the system-integration step is zero. When we run the Monte Carlo simulation for the for possible technology policies within the range of 10 to 30% capability gain, we obtain the results that are depicted in Figure 6.

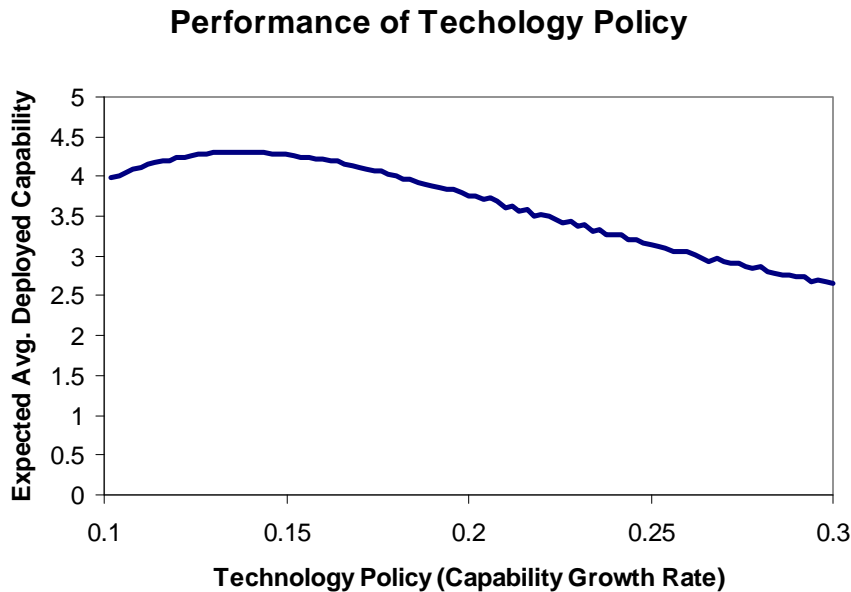


Figure 6. Performance of Technology Policies for a Single Technology

We can see in Figure 6 that there is a single optimal technology policy for our performance metric: expected average deployed capability. In fact, the optimal policy is a relatively modest 13.8% target improvement in capability for each acquisition cycle. This policy results in an expected average deployed capability of 4.31. This result seems to suggest that from a performance standpoint it is better to take smaller, more frequent steps than larger, less frequent steps. In other words, evolutionary is better than revolutionary. But is this always the case? There are two critical features of this model that we can vary. First, we can alter the integration time. In the above case, it was set to zero. But if, for example, it was set to two years, the single technology optimal policy increases to 20%. This result is reasonable because longer integration times essentially impose more overhead on the acquisition process. Consequently, it is advantageous to target a larger

increase in capability to compensate for the integration delay. However, in this model, system integration is not linked to the maturity of the technology selected. In some cases, a more immature technology may be more difficult to integrate with the rest of the system and, hence, would actually exacerbate delays.

Impact of Risk vs. Return

The second feature of the model we should consider is the shape of the curve that relates risk and return. The function used in our model is displayed in Figure 4. This curve exhibits the diminishing return to increasing risk that was mentioned earlier. But what would happen if the penalty for additional risk were more severe? In other words, what if taking on large amounts of risk resulted in very little gain in capability? As an excursion, we will assume that the relationship between the gain in capability and the upper bound of duration is determined by the following exponential relationship.

$$\text{Function 2. Capability Gain} = -0.81104 e^{-0.7 \text{Upper Bound}} + 0.300001$$

We find that under this risk-return model, the optimal single technology policy increases to 26%. If, on the other hand, we removed the diminishing returns to risk entirely, we would use the following linear relationship:

$$\text{Function 3. Capability Gain} = 0.011111(\text{Upper Bound}) + 0.077778$$

Under this function our optimal policy is 10%, the minimum allowable. This behavior is perhaps better understood visually. Figure 7 shows all three of the curves discussed. Note that all three pass through the same maximum and minimum points, so the issue is just the shape of the curve. Notice also that the exponential curve increases sharply then flattens out. The high initial derivative means that on the lower end of the curve, one can actually gain quite a bit of capability for very little risk. But the curve quickly flattens out such that each additional gain in capability requires a huge increase in risk. Thus, there is a natural optimal point. The same is true for the baseline curve. While it is not as severe, there is essentially a natural optimal increment size. For the linear curve, the derivative is constant, so the best strategy is to minimize the size of the increment. In this extreme case in which there is no integration time, we can essentially deploy infinitesimally small increments of capability continuously. Thus, the linear case ensures that the best possible capability is available at any time. While this case would be desirable, it is certainly not realistic. Something akin to the baseline case is more reasonable because, in reality, there is usually some minimum reasonable increment that can be deployed.

Risk - Return Tradeoff




Figure 7. Comparison of Different Risk-return Relationships

It is important to note that changing the scale of the risk-return relationship will certainly change the optimal policy, but here we focused on the shape. This is because the shape of the curve is what determines how aggressive the optimal policy is within the feasible ranges of capability and risk. What we can conclude from this analysis is that, from a performance standpoint, there is a natural optimal technology policy, and, except in extreme circumstances, that policy is not going to be the maximum achievable leap in capability.

Impact of Multiple Technologies

Thus, for a single technology we find that the best policy will most likely be to take small steps with more mature technologies; but what happens when a program depends on the integration of multiple critical technologies? First, we will assume that each technology provides a different capability. For example, a multi-mission surface combatant would have critical technologies that provide anti-air and anti-submarine warfare capabilities. Presumably, stakeholders for each area or capability would want to maximize their respective average deployed capability. But with multiple technologies in the same program, the actions of one affect the outcome for others. For example, the selection of an immature technology for anti-air warfare could delay the delivery of the next ship class and, consequently, delay the deployment of the next generation of anti-submarine warfare technology. From the perspective of stakeholders in anti-submarine warfare, the expected delay means that if they must wait, they should target a larger gain in capability for their area to compensate for the delay. But since program completion depends on both technologies, the reciprocating decision could actually exacerbate delays further. In order to understand stakeholder behavior when a program incorporates multiple critical technologies, we will

employ game theory. (For an introductory treatment of game theory see Gibbons (1992). For a more advanced treatment see Fudenberg and Tirole (1991).)

Game theory allows us to consider the strategies of rational competing parties. A technology policy would be the targeted percent increase for each capability for each acquisition cycle. For example, for anti-air we might target a 15% increase per cycle, while for anti-submarine we might target a 10% increase. Presumably, stakeholders for each area want to maximize the average deployed capability for their area of concern. To employ game theory, we must find the best response functions for the stakeholders for each of the capability areas. We can accomplish this by finding the optimal response to each possible action by the other player. Any intersection points between the best response functions constitute Nash equilibria. A Nash equilibrium is a stable point in strategy at which either player would be worse off if they deviated from that strategy. To demonstrate this concept, we will assume that there are two critical technologies in each acquisition program. Both have the identical risk-return behavior from the baseline case above. The resulting best response functions can be seen in Figure 8.




Figure 8. Stakeholder Best Response Functions for a System with Two Critical Technologies
(The intersection of the two functions is the Nash equilibrium.)

The plotted points in Figure 8 represent the best responses over the selected policies. Since Monte Carlo simulation was used, there is some statistical noise in these results. Consequently, a linear function was fit to the best response data for each player. We can see from the best response functions that the two players engage in reciprocating competition. That is, as each player increases his targeted capability, the best response of

the other player is to increase his as well. Since we assumed two identical players, the Nash equilibrium is symmetric and much more aggressive than what is optimal for a system with single critical technology. In fact, the equilibrium solution is for each player to target a 23% increase in capability for each acquisition cycle—resulting in an average deployed capability of 2.7 for each. This is far below the optimal single technology result of 4.31. The practical implication is that older generations of technology stay in the field much longer.

The resulting Nash equilibrium would seem to corroborate the behavior described previously. If one player chooses a particular technology policy, it is in the best interest of the other player to choose one that is just slightly more aggressive. Consequently, the first player might as well choose a more aggressive one himself, and so on. The result is an equilibrium state with a much more aggressive technology policy than we would expect from the single technology analysis. To better understand this result, let us consider the case in which there are still two critical technologies, but the two players cooperate in selecting a technology policy.

To find the best cooperative technology policies, we can search over a grid of possible policy combinations. The results are plotted in Figure 9. The plotted points form the space of all possible policy outcomes. Since we would like to maximize the performance over each capability, we must find the Pareto optimal set of policies. A Pareto optimal policy is defined such that to improve performance of one capability would mean sacrificing performance on another. The Pareto optimal set is designated by the squares in Figure 9.

Capability Performance Space

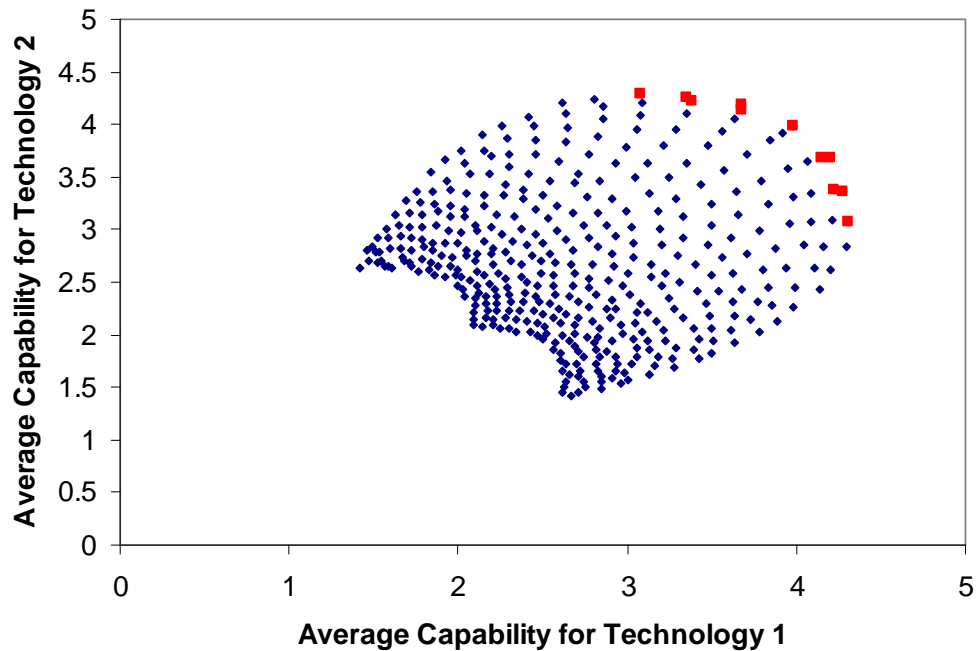


Figure 9. Performance Space of all Possible Technology Policies
(The squares indicate the Pareto optimal set of policies.)

We note that the Pareto optimal frontier allows us to trade off some performance between the two capabilities; but even so, the entire frontier is superior to the Nash equilibrium achieved in the non-cooperative case. For the sake of comparison, let us consider the optimal symmetric policy from the frontier. Under this policy, the target is a 12% improvement per acquisition cycle and an expected average deployed capability of 3.99. Note that this is a significant improvement over the non-cooperative solution but not as good as the single technology solution.

What does this result tell us? We can conclude several things. First, the best solution for a program that relies on multiple critical technologies is a cooperative one in which a small amount of capability is sacrificed from each area to bring the overall cycle-time down. We can see this sacrifice when we compare the optimal symmetric policy to the single technology policy. Thus, we see that there is a price to pay for including multiple capabilities in a single system. While there are likely cost advantages, there will be some sacrifice in performance (barring synergistic effects) because the integration of multiple technologies increases acquisition cycle-time. More importantly, however, is that the optimal solution is not stable in that it is not a Nash equilibrium. Therefore, there is always an incentive to deviate. Let us say, for example, that we select the optimal symmetric technology policy for our system. Assuming that everyone else follows this policy, it is in the best interest of anyone supporting a particular capability to push for a slightly more aggressive technology for his area. He will end up better off. But since all have an incentive to deviate, if one deviates, all will likely deviate, and we end up at the Nash equilibrium. This is exactly where we do not want to be.

To better elaborate on this point, it is instructive to consider the cartel problem from economics. In a cartel, several firms make a price-fixing agreement so that they can all earn greater profits than if they competed. Thus, they set a price higher than the market equilibrium price. However, there is an incentive to deviate. If one firm in the cartel charges slightly less than the agreed-upon price, it will capture the market and make much more money than it would by following the cartel agreement. Consequently, cartels tend to be unstable without strict monitoring and enforcement.

We see that our situation here is quite analogous. For a given system, it is in the best interest of all stakeholders and decision-makers to sacrifice a little bit of capability in each critical area in order to pursue an evolutionary rather than a revolutionary policy. However, it is always in the best interest for any given stakeholder to push for just a little bit more capability in his respective area. Thus, the best solution is unstable. This phenomenon could explain, at least in part, why the acquisition system in the Department of Defense consistently pursues revolutionary rather than evolutionary acquisition programs despite policies to the contrary. The above game theory analysis indicates that in the absence of enforcement, the rational actions of decision-makers with good intentions will lead to poor acquisition policy. The implication here is that if the Department of Defense is serious about evolutionary acquisition, it cannot expect voluntary compliance. Compliance must be enforced.

In the interests of robustness, we should consider the sensitivity of this result. If we increase the integration time to two years, the competitive policy is a 29% increase in capability per cycle for an average deployed capability of 1.78—whereas the optimal symmetric cooperative policy is an 18% increase in capability per cycle with an average deployed capability of 2.05. Thus, an increase in the integration delay makes both policies more aggressive, but the relationship between competition and cooperation is preserved.



As for the shape of the risk-return curve, if we examine the exponential case, we find that the competitive-cooperative relationship is still preserved but becomes less dramatic. The competitive policy is a 27% increase in capability per cycle resulting in an average deployed capability of 11.7—whereas the optimal symmetric cooperative policy is a 24% increase in capability per cycle with an average deployed capability of 12.74. (Note that these average deployed capabilities are very high because the exponential curve allows for large increases in capability very quickly.) Finally, what happens as we increase the number of critical technologies? It turns out that the situation gets worse. If there are three critical technologies, the competitive policy is to pursue a 33% increase in capability per cycle for an average deployed capability of 1.76. Meanwhile, the optimal symmetric cooperative policy for three identical technologies is to target a less-than-11% gain in capability per cycle resulting in an average deployed capability of 3.93.

Thus, the key result from this analysis is that when an acquisition program relies on more than one critical technology, the relationship between competitive and cooperative behavior is fairly robust. The cooperative policy yields superior performance through smaller capability increments but is unstable. Without enforcement, the situation devolves to a suboptimal Nash equilibrium that achieves inferior performance through larger capability increments.

Cost Considerations

Up until this point we have only considered performance, and we have omitted any discussion of cost. An evolutionary approach to acquisition may achieve a higher deployed performance on average, but is it more cost-effective? This question is a little more difficult to answer; it depends in large part on the relative costs to produce and deploy the replacement system (or upgrade the old system) at the end of each cycle, as well as on the relationship between technology maturity and development cost. All else being equal, we can say that there is a tradeoff between cost and performance in terms of cycle-time. More frequent, shorter cycles mean that overhead costs associated with an acquisition cycle are incurred more often. Consequently, costs will increase when the cycle-time is shorter. The relative magnitudes of cycle costs versus development costs will dictate the severity this tradeoff. More expensive development costs reduce the contribution of cycle costs as a percentage of the overall acquisition bill. Thus, the tradeoff becomes less severe. If, on the other hand, cycle costs are very high (e.g., from high manufacturing costs), increasing the length of acquisition cycles may be more appealing.

The missing piece here is the impact of technological maturity on cost. Does the inclusion of immature technology in an acquisition program require the use of more expensive development methods than if the same technology were pursued in a research and development setting? Does the inclusion of immature technology make system integration more difficult and expensive than when mature technology is employed? Conventional wisdom would suggest that the answer to both of these questions is yes, and if so, there could be an optimal technology policy that minimizes cost.

We can demonstrate, at least in a simplistic way, that it is possible to achieve lower costs through an evolutionary strategy. Let us assume that we have a system with a single capability, and we can upgrade that capability through either one large leap or multiple small steps. Both achieve the same end capability, but increasing the number of cycles to achieve it increases the maturity of the technology we use in each cycle. For example, we could achieve a 25% increase in capability all at once or through two steps that sequentially



increase capability by 11.8% each time. The capability outcome is the same in either case, but the time to achieve it may differ. In fact, if we use the baseline risk-return relationship described earlier, the one-step strategy is expected to take 5.27 years (assuming no integration time), whereas the two-step strategy is expected to take 4.46 years. Of course, the addition of integration time could erode the time advantage provided by multiple steps, but at least notionally we can see that there could be some cost advantage to taking multiple, less-risky steps. If we assume that the development costs are the same in both cases, the cost difference comes down to the overhead associated with each acquisition cycle. Again, for the sake of simplicity, if we assume that the cycle costs are the same regardless of the aggressiveness of the technology policy, we can determine the conditions under which the evolutionary strategy is more cost-effective. To make this explicit, let us define some variables.

n = the number of steps in the evolutionary policy

$D(n)$ = the expected length of each step when there are n steps

C_D = the cost rate for development work

C_O = the overhead cost for each acquisition cycle

As above, we assume that we have two policy options that achieve the same increase in capability. However, the first policy option achieves it in one step, while the second achieves it in n steps. We want to find the conditions under which the n -step policy is more cost-effective than the one-step policy.

$$\text{Equation 1. } nD(n)C_D + nC_O < D(1)C_D + C_O$$

Rearranging terms yields:

$$\text{Equation 2. } \frac{C_O}{C_D} < \frac{(D(1) - nD(n))}{n - 1}$$

Thus, we can characterize the cost-effectiveness of the evolutionary policy in terms of the ratio of the cycle costs to the development costs. For the example discussed above, the two-stage policy is more cost-effective when the ratio of the cycle cost to the development cost rate is less than 0.81. However, this analysis is admittedly oversimplified and only reveals the possibility that an evolutionary strategy could be less expensive.

Additional work is needed to model cost in a more realistic manner. Empirical studies regarding software development reveal that for large projects, the cost savings from a spiral approach can be substantial (Boehm, 2007). This occurs because the upfront investment in better systems engineering and spiral risk reduction is outweighed by the reduction in rework. This is especially true for complex software systems. This result, while particular to software systems, is suggestive and may imply that the adverse cost impacts of including immature technologies in a single acquisition cycle may be substantial. Thus, to effectively model cost, we must consider the possible approaches to developing technologies.



One approach is to leave technologies in the R&D process longer so that they are more mature when they are finally included in an acquisition program. The advantage to this approach is that technologies can be managed in a portfolio setting. That means funding can be balanced and allocated to maximize the technological options available to acquisition programs. If, for example, in the course of development, a particular technology proves to be problematic or not as effective as anticipated, funding may be shifted to an alternate approach to provide a needed capability. In contrast, once technologies are included in an acquisition program, some of this flexibility is lost. There is a great deal of commitment to a particular design approach, and it may be difficult or prohibitively expensive to change it in the event that a selected technology underperforms. Thus, to really model the cost implications of evolutionary acquisition, one would need to model the cost impacts of withdrawing technologies from the R&D portfolios at various levels of maturity. This must be relegated to future work.

Conclusions

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 increased emphasis on multi-mission or multi-capability platforms may lead to overall 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 to 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. 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.

References

- Boehm, B. (2007, April 2). [Personal communication with authors].
- Boudreau, M. W. (2006). Acoustic rapid COTS insertion—Case study. In *Proceedings of the third annual acquisition research symposium* (pp. 185-186). Monterey, CA: Naval Postgraduate School.
- DoD. (2003a, May 12). *DoD directive 5000.1*. Washington, DC: author.
- DoD. (2003b, May 12). *DoD instruction 5000.2*. Washington, DC: author.
- DoD. (2006). *Defense acquisition guidebook*. Retrieved March 16, 2006, from <http://akss.dau.mil/dag/welcome.asp>.
- Fudenberg, D., & Tirole, J. (1991). *Game theory*. Cambridge: The MIT Press.



- GAO. (2003, November 10). *Defense acquisitions: DoD's revised policy emphasizes best practices, but more controls are needed* (GAO-04-53). Washington, DC: United States Government Accountability Office.
- GAO. (2006, April 5). *Defense acquisitions: Actions needed to get better results on weapons systems investments* (GAO-06-585T). Washington, DC: United States Government Accountability Office.
- GAO. (2006, April 13). *Defense acquisitions: Major weapon systems continue to experience cost and schedule problems under DoD's revised policy* (GAO-06-368). Washington, DC: United States Government Accountability Office.
- GAO. (2006, December 21). *Defense contracting—Questions for the record* (GAO-07-217R). Washington, DC: United States Government Accountability Office.
- Gibbons, R. (1992). *Game theory for applied economists*. Princeton: Princeton University Press.
- Johnson, W. M., & Johnson, C. O. (2002). The promise and perils of spiral acquisition: A practical approach to evolutionary acquisition. *Acquisition Review Quarterly*, Summer, 175-188.
- Pennock, M.J., Rouse, W.B., & Kollar, D. L. (2007). Transforming the acquisition enterprise: A framework for analysis and a case study of ship acquisition. *Systems Engineering*, 10(2).
- Smaling, R., & de Weck, O. (2007). Assessing risks and opportunities of technology infusion in system design. *Systems Engineering*, 10(1), 1-25.



2003 - 2006 Sponsored Acquisition Research Topics

Acquisition Management

- Software Requirements for OA
- Managing Services Supply Chain
- Acquiring Combat Capability via Public-Private Partnerships (PPPs)
- Knowledge Value Added (KVA) + Real Options (RO) Applied to Shipyard Planning Processes
- Portfolio Optimization via KVA + RO
- MOSA Contracting Implications
- Strategy for Defense Acquisition Research
- Spiral Development
- BCA: Contractor vs. Organic Growth

Contract Management

- USAF IT Commodity Council
- Contractors in 21st Century Combat Zone
- Joint Contingency Contracting
- Navy Contract Writing Guide
- Commodity Sourcing Strategies
- Past Performance in Source Selection
- USMC Contingency Contracting
- Transforming DoD Contract Closeout
- Model for Optimizing Contingency Contracting Planning and Execution

Financial Management

- PPPs and Government Financing
- Energy Saving Contracts/DoD Mobile Assets
- Capital Budgeting for DoD
- Financing DoD Budget via PPPs
- ROI of Information Warfare Systems
- Acquisitions via leasing: MPS case
- Special Termination Liability in MDAPs

Logistics Management

- R-TOC Aegis Microwave Power Tubes
- Privatization-NOSL/NAWCI
- Army LOG MOD
- PBL (4)



- Contractors Supporting Military Operations
- RFID (4)
- Strategic Sourcing
- ASDS Product Support Analysis
- Analysis of LAV Depot Maintenance
- Diffusion/Variability on Vendor Performance Evaluation
- Optimizing CIWS Lifecycle Support (LCS)

Program Management

- Building Collaborative Capacity
- Knowledge, Responsibilities and Decision Rights in MDAPs
- KVA Applied to Aegis and SSDS
- Business Process Reengineering (BPR) for LCS Mission Module Acquisition
- Terminating Your Own Program
- Collaborative IT Tools Leveraging Competence

A complete listing and electronic copies of published research within the Acquisition Research Program are available on our website: www.acquisitionresearch.org





Acquisition research Program
Graduate school of business & public policy
Naval postgraduate school
555 DYER ROAD, INGERSOLL HALL
MONTEREY, CALIFORNIA 93943

www.acquisitionresearch.org