ARROYO CENTER

The RAND Corporation is a nonprofit institution that helps improve policy and decisionmaking through research and analysis.

This electronic document was made available from www.rand.org as a public service of the RAND Corporation.

Skip all front matter:

Support RAND

Purchase this document

Browse Reports & Bookstore

Make a charitable contribution

For More Information

Visit RAND at www.rand.org

Explore the RAND Arroyo Center

View document details

Toward Affordable Systems II

Portfolio Management for Army
Science and Technology Programs
Under Uncertainties

Brian G. Chow, Richard Silberglitt, Scott Hiromoto,
Caroline Reilly, Christina Panis

RAND ARROYO CENTER

Summary

S&T plays a central role in the ability of the Department of Defense (DoD) to field the advanced weapons systems that give the U.S. military its unmatched technological superiority. In view of this fact, DoD's program for acquiring new systems has long been linked with its S&T programs in basic and applied research and advanced technology development.[1] In 2008, for the first time, DoD explicitly documented a process for technologies coming out of its S&T programs to be integrated into new systems at every stage of acquisition (see Figure S.1). Under previous policies, the linkage between S&T and acquisition typically occurred at the initial stage of the acquisition process, prior to Milestone A, while the basic system concept was being refined and before the total costs to develop, field, and operate a new system were assessed. These lifecycle costs were part of an analysis of alternatives (AoA) due at Milestone A in the acquisition process (shown in Figure S.1), just before a new system entered the "technology development" stage. Consequently, the military's S&T planners were not called upon to consider the lifecycle costs of the systems in which new technologies were being used until an AoA was required. They had the costs of their own S&T programs to manage but could do so more or less independently of total systems costs. The 2008 policy reflects a department-wide emphasis on technology insertion at every stage of the acquisition process. To support such insertions, planners would need to make lifecycle cost estimates at the completion of the S&T programs, since the next step, the AoA analysis, where lifecycle cost estimates are typically made, may be skipped.

A second significant change in the acquisition policy environment occurred in 2009 when Congress established a DoD position, the director of Independent Cost Assessment. The director is tasked with conducting independent estimates of the cost of new major military systems. For any new system, this director needs to consider "trade-offs among cost, schedule, and performance objectives" (P.L. 111-23, p. S. 454-1).

---

[1] Advanced technology development "includes development of subsystems and components and efforts to integrate subsystems and components into system prototypes for field experiments and/or tests in a simulated environment." See Office of the Under Secretary of Defense, Comptroller, *DoD Financial Management Regulation, Volume 2B: Budget Formulation and Presentation*, "Research, Development and Evaluation Appropriations," Chapter 5 (July 2008), July 6, 2000, p. 5-2.

Figure S.1
Links Between Army S&T Programs and the Defense Acquisition Management System
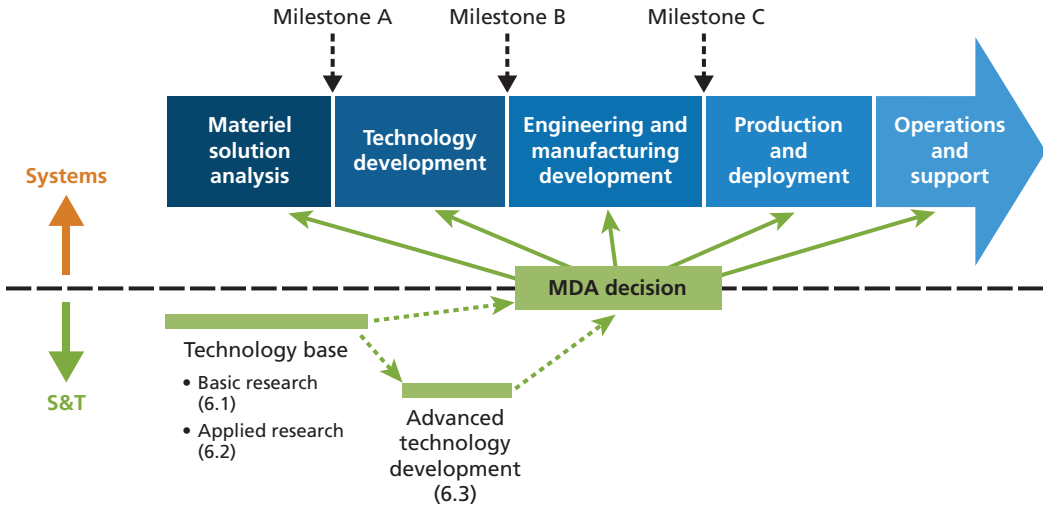
SOURCE: Simplified graph from Brown, 2008.
RAND *MG979-S.1*

Both of these policy developments highlight a rising concern with the cost of advanced weapons systems. Therefore the Army, as well as the other services, will need to rely more than ever on its scientists and engineers to design systems that will not only provide necessary capabilities, but that can also be acquired and operated affordably. Decisions made in the S&T stages of the acquisition process are vital in determining the total lifecycle costs of a new system. Historically, 70 percent of the overall cost of a system is incurred once it enters the final Operations and Support phase (U.S. Army, 2006, pp. 26–27). It is much easier to take steps to reduce those late-stage costs early in the acquisition process, before a design is fully developed and hardened. In this way, performance and lifecycle costs can be weighed against each other to determine what trade-offs may be viable.

The Army faces continual challenges in selecting portfolios of S&T projects that will meet future capability needs at an affordable overall cost. Uncertainty about whether all funded S&T projects will succeed compounds these challenges, as does the ever-present possibility that changes in the economic or strategic environment will compel some capability requirements to be altered.

A Process to Make Early-Stage Decisions That Will Lead to Effective and Affordable Systems

With this need in mind, the RAND Corporation developed a process to help the Army incorporate lifecycle cost into S&T planning under uncertainty.[2] This process offers several advantages. It equips the Army to consider lifecycle costs at the S&T stage so that the Army can design more cost-effective systems. It can complement the Army's existing process for managing its S&T programs. Moreover, the process will create a new opportunity for dialogue among stakeholders and allow different viewpoints and perspectives to be analyzed objectively in the process of building an S&T portfolio.

The Army's S&T managers can use the RAND process as an aid in carrying out two fundamental planning tasks:

- **Mapping supply and demand:** The process first broadly identifies where the Army may encounter problems meeting requirements (the demand) with existing S&T projects (the supply). It then refines this broad map of supply and demand, taking into account that some S&T projects in the portfolio will not lead to a fielded system, because either the project fails or the system is not as cost-effective as another.
- **Selecting an optimal S&T portfolio:** The process next identifies how the Army's S&T planners can manage expected or unexpected budget cuts, while keeping chances of meeting requirements as high as possible.

Mapping Supply and Demand

With any S&T portfolio, inevitably certain projects will succeed, while others will fail. (For the purposes of this study, we consider a successful S&T project one that meets its system performance and cost objectives and is suitable for further development into a fielded weapons system.) Failure to take this reality into account during planning can create a false impression of how well a given S&T portfolio will meet the Army's capability needs. An accurate map of mismatches in supply and demand is the product of two screening steps, one of which accounts for uncertainty. In the first step, the RAND process helps planners establish a baseline for measuring how well supply meets demand by assuming that all S&T projects in a given portfolio will succeed. But, potentially, enough projects could fail so that all requirements cannot be met. Consequently, in a second step, the process enables planners to analyze how this uncertainty affects that portfolio's ability to meet requirements. Through this process,

---

[2]   The process is presented in two monographs. The first—Chow, Silberglitt, and Hiromoto, 2009 (hereafter referred to as TAS-1)—provides the basic model. The current study (TAS-2) adds a simulation model to deal with uncertainty.

planners can clearly see which areas of demand are at risk of not being met and where the portfolio may need to be adjusted accordingly.

Step One: Broadly identify where the Army may encounter problems meeting requirements (demand) with existing S&T projects (supply). A comprehensive picture of supply and demand takes the form of a matrix (see Figure S.2). For the demand in our demonstration of the RAND process, we use the force operating capability (FOC)–gap requirements defined by 2006 data from TRADOC/ARCIC.[3] These requirements are displayed from left to right along the top of the matrix. For this demonstration, the Army Technology Objectives (ATOs) constitute the supply. We use these because they are the Army's highest-priority S&T projects and are examples of S&T projects that can lead to specific systems.[4] The ATOs are displayed from top to bottom along the left side of the matrix.[5] Each row in the matrix displays how many capability gaps a given ATO addresses within each FOC requirement. The columns indicate how many ATOs address that gap.

Drawing on this matrix and assuming that none of the ATOs in this portfolio will fail, we can identify the overall performance of the current set of ATOs and expose potential problem areas—that is, where capability gaps are at risk of not being met (see Figure S.3). Under these ideal conditions, where 100-percent success is a given, the overall set of ATOs in the pipeline can satisfy nearly all of the Army's current capability needs: Ten out of the 11 meet at least 100 percent of the FOC requirement. Only one problem area emerges clearly: Existing ATOs meet only 57 percent of requirement 10, "Training, Leadership, and Education."

However, more-detailed consideration of this map provides important information about where the Army's S&T planners might want to adjust the roster of ATOs in the existing portfolio. Two of the requirements—4 and 6—are only just met. Both constitute red flags, where the current set of ATOs is at risk of not being able to meet these needs should certain projects fail. These are areas where new ATOs may be needed. On the other hand, requirements 3, 5, and 8 are overmet, indicating redundancy among projects in the portfolio. Here may lie opportunities to reduce funding of some existing ATOs that are providing too high a level of redundancy. The RAND process can show how to save money by terminating some ATOs and using that money to fund some new ATOs targeted at requirements that will be inadequately met by existing ATOs, resulting in a higher chance of meeting all requirements at the same or even lower total cost.

---

[3]   To allow for analysis according to FOC, RAND assigned each of the TRADOC/ARCIC–defined capability gaps to the FOC to which it best applied. In this monograph, the capability gaps are listed only by FOC and number. The Excel spreadsheet that identifies the gaps came from an unpublished TRADOC/ARCIC report on force-capability-gap analysis using 2006 data.

[4]   Any desired set of S&T projects expected to lead to specific systems could be analyzed using the RAND process.

[5]   The ATOs are listed by technology type and title.

Figure S.2
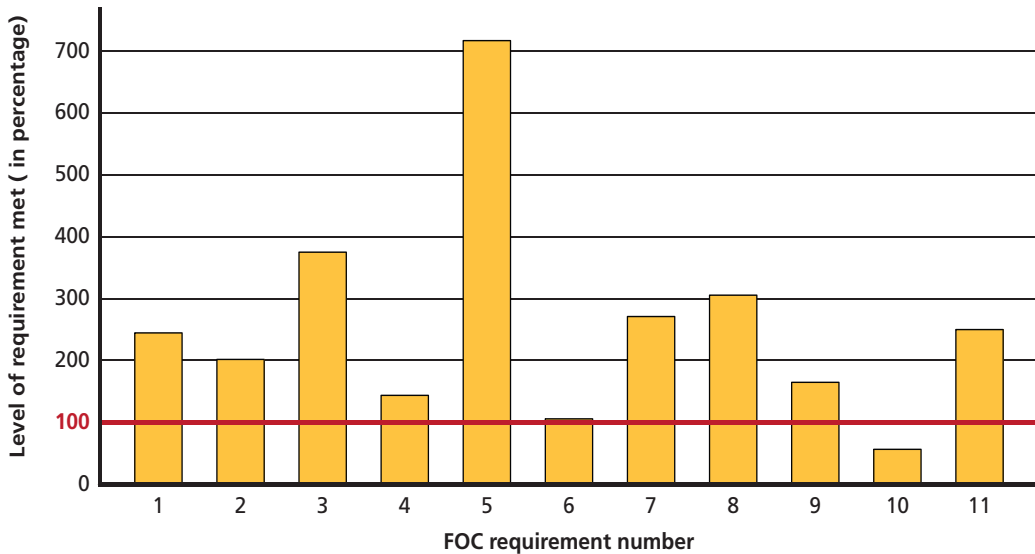**Matrix of Supply and Demand (Gap-Space Matrix)**
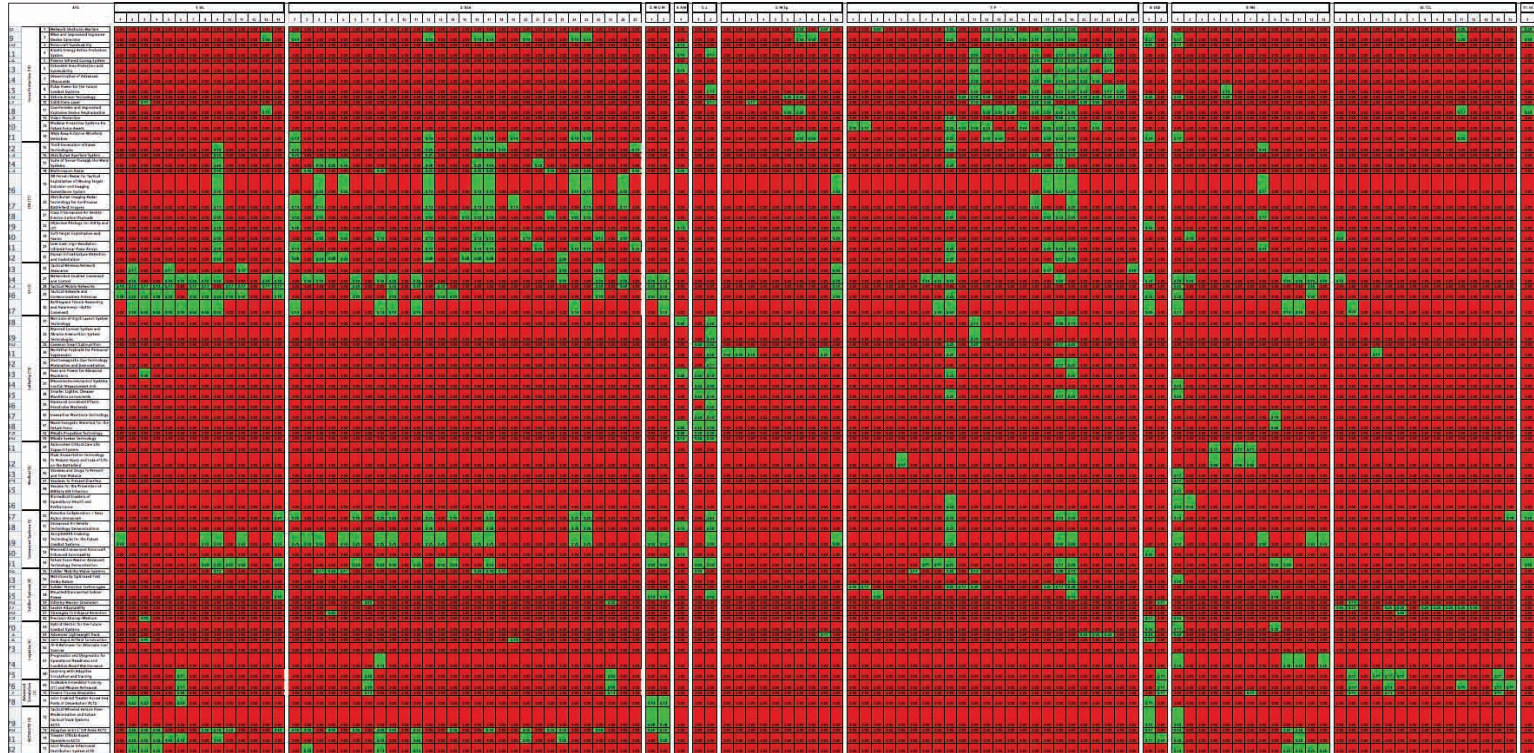
NOTE: A larger version of this matrix is shown in Chapter Three, Figures 3.1–3.8.

Figure S.3
Map of the Match Between Supply and Demand, Assuming a 100-Percent Success Rate for All ATOs in the Existing Portfolio

NOTE: We obtained these estimated percentages by adding the independent contributions of each ATO to each capability gap using the gap-space method defined in Chapter Two, as described in Chapter Three.
RAND *MG979-S.3*

This initial map of supply and demand, which assumes a 100-percent rate of success for all ATOs, is a solid starting point for planning efforts. It shows the Army's S&T managers the overall performance of existing ATOs and gives them a broad idea of where they might divert funds toward new ATOs that target unmet requirements or those at risk. The levels of each FOC requirement met—that is, the 11 percentages displayed from left to right in the bar graph—constitute a baseline that planners can use to refine the map in the next step of the process.

Step Two: Refine this map of supply and demand, given that some S&T projects in the portfolio will not lead to a fielded system. Assuming a 100-percent rate of success for all ATOs, although informative, is ultimately a theoretical exercise: Some S&T projects will inevitably fail. Because of this, the Army cannot be sure that existing projects will meet all of its capability gaps. When the success of S&T projects is uncertain, planners must unavoidably deal with probability. Our concept of a *feasible percentage* is a way of managing the S&T portfolio under uncertainty. The feasible percentage indicates the probability that a given portfolio of S&T projects will meet a defined requirement level for each FOC, with the possibility built in that some projects in the portfolio will fail—however, without knowing in advance which ones these will be.

The RAND process takes account of uncertainty by employing a model that uses as input the probability of success of each S&T project in the existing portfolio. To

demonstrate this model, we run a computer simulation 10,000 times that follows three principal steps:

- Assume a 10-percent probability of failure for all ATOs.[6]
- Allow projects in the existing portfolio to fail randomly at that rate.
- Estimate the effects.

The feasible percentage is the percentage of times out of those 10,000 runs that the portfolio meets a defined requirement level for each capability gap.

In our demonstration of the RAND process, when ATOs in the existing portfolio fail randomly at a rate of 10 percent, as a set, they fall far short of meeting the Army's baseline percentage levels. Recall that the baseline is the outcome of Step One in the process of mapping supply and demand: namely, the percentage of each FOC requirement met by the current portfolio of ATOs *when all projects in that portfolio succeed*. Without the possibility of any ATOs failing, the overall match between supply and demand is very good; the existing portfolio can meet 100 percent of all the requirements except for 10, which is met at 57 percent. But with the 10-percent failure rate built in, the chance that the current portfolio will meet these baseline percentages—what we call the *feasible percentage*—drops sharply. Even with a 90-percent rate that any given project will succeed, the existing portfolio of ATOs will likely have only a 16-percent chance of matching the baseline FOC percentages (100 percent for each FOC requirement except 57 percent for requirement 10). At this rate of failure, the existing ATO portfolio would clearly not be meeting all of the percentages in the baseline. But it is unclear how much each percentage in the baseline needs to be reduced in order for the portfolio to have a good chance of meeting these lowered requirements.

Consequently, the Army needs a way to refine the initial map of supply and demand when the possibility of failure is introduced. The RAND process includes such a routine. Within each FOC requirement, a number of ATOs will contribute to it. For illustration, assume that these ATO contributions together amount to 110 percent of a specific FOC requirement. The routine starts by asking what if the ATO making the largest contribution, say 30 percent, failed to be completed. Then, the rest could at best meet 80 percent (i.e., 110 − 30 percent) of that FOC requirement. Critical to the routine is a different perspective: If that requirement were to be reduced to 80 percent, planners could afford a failure in any one of the ATOs that make a contribution to that requirement, since the 80-percent figure was arrived at assuming the *largest* contributor failed. Indeed, for the simulation (with the largest contributor to each FOC assumed to fail) with reduced requirements (to 88 percent in FOC 6, to 36 percent in FOC 10, but no reduction in others) shown in row two of Table S.1, the feasible percentage jumps

---

[6]   The model allows different probabilities of failure for different projects. However, because planners generally do not know which projects will actually fail, using the historic failure rate for every project is a reasonable approximation. For this demonstration, we assume a single 10-percent failure rate.

from 16 to 73. The last row in the table shows that the feasible percentage increases to 99.8 if each FOC requirement is reduced to a level that can accommodate the failure of the four largest contributing ATOs. In other words, a simultaneous failure of any four ATOs within each category would still allow the set of reduced requirements to be met. The price to pay to drastically increase the feasible percentage from 16 to 99.8 is the need to reduce the requirements from 100 to 75 percent, 58 and 50 percent for FOC 4, 6, and 11, respectively; and FOC 10 from 57 percent down to 23 percent. This is an attractive price because we can design new ATOs targeted toward these requirement gaps to bring the requirements back to the baseline level (row one). The cost of doing so is much lower than that of adding new ATOs with a targeting pattern similar to the existing ATOs in order to bring the feasible percentage from 16 to 99.8.

Selecting an Optimal S&T Portfolio

Once the RAND process has been used to precisely determine which parts of the supply the existing ATOs should meet, with the rest to be met by new ATOs, it can then help the Army select an optimal S&T portfolio among the existing ATOs. In other words, only certain existing ATOs should be continued, and the money saved from the termination of the rest is more cost-effectively spent on new ATOs. The process also provides S&T planners a means of identifying how to manage (expected or) unexpected budget cuts, while keeping the chances of meeting requirements as high as possible. Budget cuts may be made to the total budget for developing the new systems and getting them through procurement, then fielding, operating, maintaining, and, finally, decommissioning them—the total implementation budget—or to the budget for funding ATOs to the end of the S&T phases—the total remaining S&T budget. Changes to either budget will alter the probability that ATOs will cover requirements—in other words, the feasible percentage. The sum of these two budgets is called the total remaining lifecycle budget.

The map of supply and demand deals purely with the ability of a portfolio to meet requirements. But in the real world of acquisitions, cost is an ever present concern. When costs are factored in, the goal of S&T planners becomes the *sweet spot*

Table S.1
Refined Map of the Match Between Supply and Demand

FOC Requirement Number and Percentage											Feasible Percentage
1	2	3	4	5	6	7	8	9	10	11	
100	100	100	100	100	100	100	100	100	57	100	16
100	100	100	100	100	88	100	100	100	36	100	73
100	100	100	75	100	58	100	100	100	23	50	99.8

NOTE: The final row is the new baseline.

between the probability of meeting requirements and the affordability of doing so. Such a sweet spot indicates the total remaining S&T budget and the total remaining lifecycle budget that should be spent on the selected existing projects and their systems and the feasible percentage that will result. If a planner wants a higher feasible percentage, the sweet spot tells us that it is less costly to fund new projects, as opposed to funding more existing projects. To find this sweet spot for existing projects, the S&T planner must make trade-offs between the feasible percentage and affordability. Our process shows us these trade-offs.

To explore what happens to the feasible percentage when budgets are reduced, the RAND process uses a linear programming model and simulation, as described in Chapters Two and Four. In our demonstration, we first adjust the baseline percentage levels for each FOC requirement to match the final map of supply and demand (i.e., the bottom row of Table S.1) and then assume a total remaining S&T budget of $3.1 billion, enough to allow funding of all the ATOs to completion of S&T (see Figure S.4). Under these conditions, as the figure shows, cutting the total remaining lifecycle budget nearly in half—from $138 billion[7] to $67 billion—gives the seemingly surprising result that the large cut has practically no effect on the feasible percentage: The

Figure S.4
The Effect of Cuts in the Total Lifecycle Budget on the Probability That Existing ATOs Will Meet Requirements
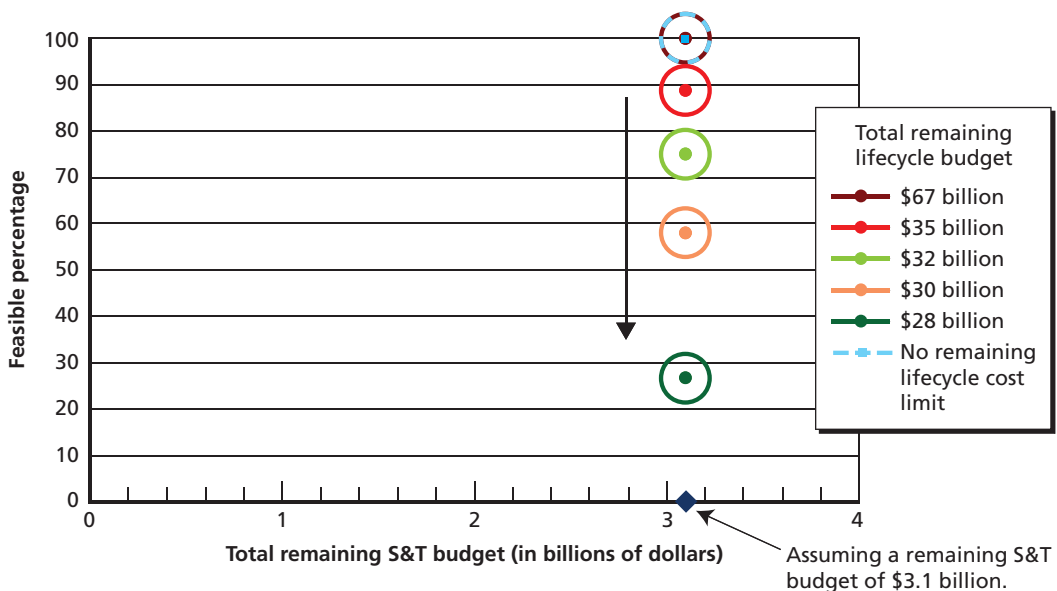
NOTE: The feasible percentage is based on an adjusted baseline requirement.

---

[7]   This would be the total lifecycle cost if all systems resulting from the 75 existing ATOs were developed and fielded, designated "no total remaining lifecycle cost limit."

probability that existing ATOs will meet the requirements remains the same.[8] Cutting the lifecycle budget in half yet again only brings a small loss in feasible percentage. When the budget drops from $67 billion to $35 billion, the probability of meeting requirements becomes 88 percent. This is a big gain in affordability for only a small trade-off in probability. But the loss is not linear: Trimming only $5 billion more makes probability suddenly drop sizably, from 88 to 58 percent. Shaving off another $2 billion spurs a decline of 31 percent more. Given these trade-offs, the most cost-effective total lifecycle budget for the systems resulting from these ATOs is $35 billion.

Similarly, the S&T budget can be cut with some flexibility only up to a certain point (see Figure S.5). When we take the most cost-effective lifecycle budget of $35 billion as an anchor point and now focus on altering the remaining S&T budget, once that budget falls below $1.5 billion, the probability of meeting requirements plummets. As the figure shows, the same pattern holds within other total lifecycle budgets as well. This cautions the S&T planner not to drop the S&T budget below this $1.5 billion threshold. For the current case, with a lifecycle cost of $35 billion, the opti-

Figure S.5
The Sweet Spot Between Performance and Affordability in the Lifecycle and S&T Budgets

---
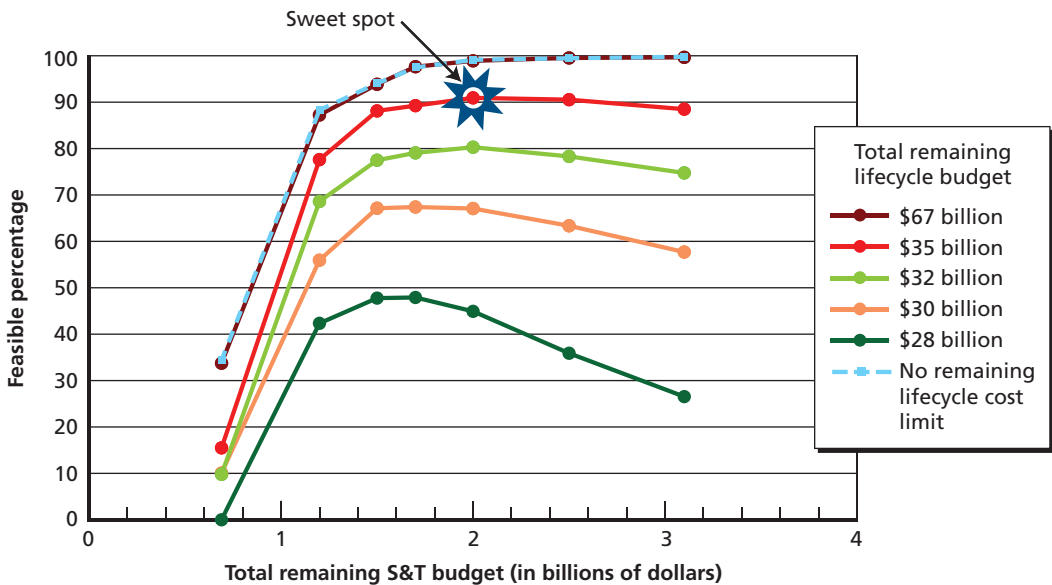
[8]   This result reflects a high level of redundancy in the systems resulting from the ATOs. Therefore, only some, and far from the complete set, of the successful ATOs are needed to be further developed into fielded systems in order to meet all requirements—and $67 billion will suffice.

mal S&T budget is $2 billion, because that amount allows the probability of meeting requirements—the feasible percentage—to reach its peak. Thus, the sweet spot for supporting the existing projects is a total remaining S&T budget of $2 billion and a total remaining lifecycle budget of $35 billion. There the Army will get a 91-percent probability (feasible percentage) of meeting its FOC requirements.

Once the lifecycle and S&T budgets are set at the sweet spot, the job then becomes to select the most cost-effective ATOs. Different portfolios of ATOs are of course possible for each combination of total and S&T budgets; various portfolios may perform better or worse. The RAND process includes an algorithm that automatically tries out different combinations of ATOs meeting any given set of budget constraints. Uncertainty is taken into account, with only successful ATOs included in the combinations. The results suggest which ATOs to keep and which to discontinue to get the highest possible feasible percentage—in other words, the most cost-effective ATOs.

To demonstrate, we take the optimal budgets from our analysis: $35 billion for lifecycle costs and $2 billion for S&T costs. Within these cost constraints, the model suggests 53 ATOs to keep and 22 to discontinue (see Figure S.6). In Figure S.6, we also compare the model selection with the best of many customary metrics we investigated. ATO 58 is rejected, even though its ratio of benefit-to-S&T cost is similar to many others selected. In contrast, three ATOs (48, 46, and 47) are selected although they have ratios even worse that those rejected ATOs with very poor ratios. These recommendations cannot be derived from a simple analysis of benefit-to-cost ratios; the model built into the RAND process is needed. It accurately represents the complex interplay between the costs of each ATO, requirements met by multiple ATOs, and the uncertainty in success of the ATOs.

Testing the RAND Process as Part of the Army's Annual Process

The RAND process could be integrated into the Army's existing annual process for making decisions about the S&T portfolio. This would be an effective way for the Army to test it, by using it to produce real-life results. The key stakeholders involved would be the S&T project managers, the Deputy Assistant Secretary of the Army for Research and Technology (DAS [R&T]), and the Warfighting Technical Council (WTC). It would take two fiscal years to run the whole test,[9] with the first year needed to establish a baseline. For the requirements, the Army could use the latest capability needs requirements provided by TRADOC/ARCIC.

---

[9]   Of course, the test can be completed much sooner if the Army commits greater resources to it.

Figure S.6
The RAND Model's Selection of Those ATOs to Keep and Those to Discontinue
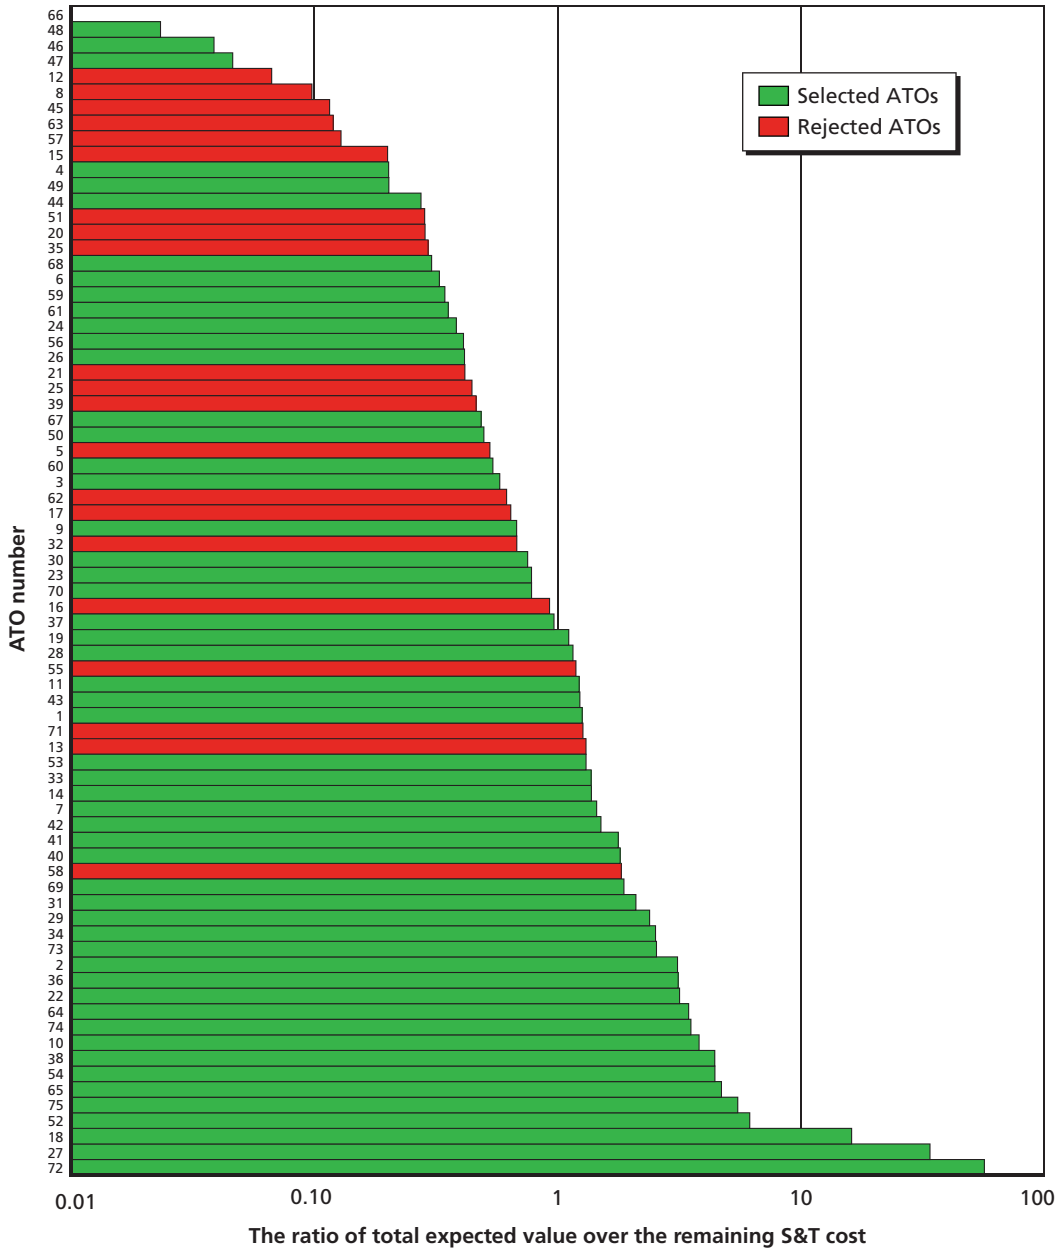
NOTE: The optimal budgets of $35 billion for lifecycle costs and $2 billion for S&T costs are used to derive the benefit-to-cost ratio.

RAND *MG979-S.6*

The First Year

In the first year, the S&T project managers and DAS (R&T) would work together to establish an initial baseline for both performance and cost. This process would follow several steps:

- S&T project managers would estimate (1) the contribution that each S&T project that leads to a fielded system will make to Army capabilities and (2) the lifecycle cost of the system derived from each S&T project.
- They would then provide these data to DAS (R&T).
- At the same time, DAS (R&T) would assemble data on the S&T cost to complete each project.
- Taking the estimates provided by the S&T project managers, DAS (R&T) would use a Delphi (or other) method to gather expert opinion on whether the estimates are appropriate, too high, or too low.
- DAS (R&T) would then apply RAND's process to inform S&T project managers of his office's current assessment of their projects. The process includes RAND's linear programming model and simulation (as described in Chapters Two and Four) and would provide outputs such as those shown in Table S.1 and Figures S.4–S.6.
- As the final step in the first fiscal year, DAS (R&T) would invite the S&T project managers to improve their baseline estimates.

The Second Year

In the second year, estimates would be refined and decisions implemented. WTC would now join the S&T managers and DAS (R&T) in the process.

- In the first six months of the year, the S&T project managers would provide their adjusted estimates and justify to WTC any differences from their baseline estimates.
- In the second six months of the year, DAS (R&T) would reapply RAND's process, with adjusted estimates where relevant, and provide the revised outputs to WTC.
- Toward the end of the year, WTC would make decisions based on the revised outputs and take any necessary corrective actions with respect to the S&T portfolio.

How Would Incorporating the RAND Process Improve the Army's Annual Process and Help It Better Manage Its S&T Portfolio?

RAND's process can help Army S&T planners monitor the expected performance of *and* the lifecycle cost of new Army weapons systems, and then weigh the trade-offs

between the two, allowing adjustments to be made at the S&T stages. These early adjustments can improve the Army's ability to build affordable new systems that also satisfy the Army's capability requirements. Integrated into the Army's existing S&T decision process, the RAND process can be a useful management tool for Army S&T planners. It can enable them to make better decisions both about individual systems and about the Army's S&T portfolio as a whole. It can secure sizable cost savings in the long-term, making it less likely, for instance, that needed weapons systems will be cancelled because they are too costly. It can help planners deal wisely and effectively with suboptimal budgets. Finally, it provides a means with which they can manage the inevitable uncertainties involved in planning by permitting unbiased "what-if" analyses of variations in the expected performance and cost of a portfolio of ATOs when certain projects fail or budgets fluctuate, as in real life they are bound to do.