Theory and methods for supporting high-level decision making

Paul K. Davis,* The RAND Corporation and Pardee RAND Graduate School
James P. Kahan,† The RAND Corporation
1776 Main Street, Santa Monica, CA 90407-2138

ABSTRACT

High-level decision makers face complex strategic issues and decision support for such individuals needs to be top-down, and to use representations natural to their level and particular styles. Decision support should focus on objectives; uncertainties, which are often both large and deep; risks; and how to do well despite the uncertainties and risks. This implies that decision support should help identify flexible, adaptive, and robust strategies (FAR strategies), not strategies tuned to particular assumptions. Decision support should also have built-in zoom capability, since decision makers sometimes need to know the underlying basis for assessments in order to review and alter assumptions, and to communicate a concern about details that encourages careful work. These requirements apply to both strategic planning (e.g., force planning in DoD or the Services) and operations planning (e.g., a commander's war planning). This paper discusses how to meet the requirements and implications for further research and enabling technology.

Keywords: Decision support, strategic planning, adaptive decision making, robust decisions, uncertainty, risk

1. INTRODUCTION

The purposes of this paper are (1) to describe a framework for supporting high-level decision makers in defining and choosing among courses of action (COAs); (2) to do so while accounting for differences in context and decision maker styles; (3) to illustrate some analytical methods and tools to provide this support; and (4) to draw implications for next steps in related research and development. The framework described can be applied both to force planning and to the early phases of operations planning. The current paper draws from two longer reports, now in preparation, which provide more context, detail, and references to the literature.[‡]

Much of what we discuss in this paper has broad applicability, but the particular paper was motivated by (1) military decisions that might be faced by, e.g., a joint task force commander or his Joint Forces Air and Space Component Commander (JFACC); and (2) decisions faced by defense planners (sometimes called "force planners"), such as the Secretary of Defense or Service chiefs. It was also developed with a specific eye towards system-of systems thinking, in which political, military, economic, social, infrastructure, and informational (PMESII) systems all play a role in understanding the decision.

The paper proceeds as follows. Section 2 discusses issues arising in higher-level decision making. Section 3 reviews aspects of decision making theory, noting distinctions between rational-analytic and naturalistic approaches. Section 4 describes a framework that moves toward such a synthesis. Section 5 presents concrete examples of new methods and tools. Finally, Section 6 discusses what is needed next from research and technology to better enable the central ideas of this paper.

[*] E-mail: pdavis@rand.org; Telephone: 310 451-6912

[†] E-mail: James_Kahan@rand.org; Telephone: 503 774-1386

[‡] Paul K. Davis and James P. Kahan, "Theory and Methods for High-Level Decision Support," RAND; and Paul K. Davis and Russell D. Shaver, "An Analytic Framework for Capability Reviews and Tradeoffs Among Options," RAND; both in preparation.

Enabling Technologies for Simulation Science X, edited by Dawn A. Trevisani
Proc. of SPIE Vol. 6227, 622701, (2006) · 0277-786X/06/$15 · doi: 10.1117/12.667026

2. GOALS, OBJECTIVES, EFFECTS, UNCERTAINTIES, AND RISKS

High-level decision makers are responsible for establishing visions and objectives, and for assuring that these are mapped into concrete objectives ("operational objectives") and actions (tasks to be accomplished). There is a long history of related methodologies, such as Strategies to Tasks.[1] The new doctrinal concept of effects-based operations (EBO) creates an intermediate step between objectives and actions by identifying the desired effects. EBO has been discussed by a number of authors[3-7] and, while appropriately quite controversial, is becoming doctrinally established.[8]

The implication for decision support (DS) is that high-level decision makers should be seeing summary materials that focus attention on the ability of options to achieve all objectives and related effects—not just a subset easy for analysts to deal with, but *all* of them. This need for comprehensiveness applies both to force planning and operations planning. Providing such a comprehensive view is part of taking a good "systems approach." A related effort is seeking to develop what has been called a Commander's Predictive Environment (CPE).[9] Here too there is a long history of related methodologies, such as that on "understanding commanders' information needs."[10]

Both EBO and CPE involve a second important class of issues, involving uncertainty and risk. High-level planners are often confronted with enormous uncertainties, *some* of which imply risks and others of which signal potential opportunities. More than a few of the uncertainties are likely to be "deep," by which we mean materially important uncertainties that cannot be adequately treated as simple random processes and that cannot realistically be resolved at the time they come into play. Such deep uncertainty must be acknowledged in the decision planning process—either by considering alternative COAs or seeking to buy time and/or information. An example of deep uncertainty is imagining the strategy of a future adversary commander in a future war in a future set of circumstances, all of which are hypothetical and unknowable.

Both of us have emphasized the uncertainty issue for some years, in both non-defense contexts,[11] and defense work that includes parametric "scenario-space" methods for confronting the uncertainties systematically.[12-14] The primary conclusion—looking across many problem domains—is that *the solution to uncertainty is to find flexible, adaptive, and robust strategies (FAR strategies).* Flexibility refers to the ability to perform different missions. Adaptiveness refers to the ability to adjust readily to diverse circumstances. Robustness refers to the ability to withstand both foreseen and unforeseen shocks. This conclusion (often referred to simply in terms of adaptiveness or robustness) has also been discussed for organizations generally.[15] The goal of finding a FAR strategy is in dramatic contrast to seeking an "optimal" strategy that assumes a particular future, as occurs when people take particular planning scenarios too seriously. Official emphasis of such considerations began to appear in DoD's first Quadrennial Defense Review (QDR) in 1997 and is at the core of DoD's capabilities-based planning (CBP), introduced in the 2001 QDR and reinforced by the 2006 version.

DoD officials are also particularly concerned about managing risks. They refer to the need to identify areas in which "more risk can be taken," but in more classic terminology the issue is actually the eternal challenge of balancing budgets with limited resources—of finding bill payers as well as new claimants. A decision support system (DSS) should address risk directly and encourage choice of FAR strategies within an economic framework.

For commanders, the resources issue is less about budgets than about people and materiel. The pinch is not always large. During the Balkan conflicts of the 1990s, the air forces were more than ample and the most difficult choices could be avoided. However, the U.S. ground commanders in Iraq are acutely aware of how limited their resources actually are and in future conflicts, such problems will again beset both air-and-space and maritime commanders, especially in situations with limited strategic warning and limited use of foreign bases.

It follows that a DSS should highlight issues of uncertainty, risk, and choice. This must be done hierarchically, at different levels of abstraction, but these issues must not be relegated to mere footnotes. Their *essence* must rise to the top. This is often not so easy, especially given the differing styles of decision makers as discussed in the next section.

3. SELECTED REVIEW OF DECISION MAKING THEORY

3.1 Background

As discussed in a recent RAND review of modern decision making theory conducted for the Air Force Research Laboratory,[16] the mindsets and methods associated with so-called rational-analytic and naturalistic methods to decision making are sharply in conflict (even if the underlying science is not). Analysts and decision-support developers tend toward their approximation of the rational-analytic style of comparing and choosing among multiple options, whereas many real-world high-level leaders are oriented toward intuitive approaches based in part of their experiences, and on taking prompt action (with adjustments later as necessary). President Bush describes himself as very much intuitively inclined, as do many other leaders of nations, armies, and corporations. MBA students are often exhorted to worry most about getting to market quickly, with the admonition not to be paralyzed by analysis.

For many years, it was widely believed that the rational-analytic style was obviously "right" (although limited in applicability because of "bounded rationality") and that the task was therefore to help decision makers adhere to it, rather than to fall prey to the many well known cognitive biases. In more recent years, however, it has been noted that the intuitive style of decision making is often very effective and should therefore not be scorned. Further, it has increasingly been appreciated that efforts to implement rational-analytic methods often do violence to many of the real issues, much as overly "rational" analysis in our day-to-day world often lacks common sense. Analyses may, for example: (1) suppress important "soft" factors; (2) ignore considerations of personality and political context; (3) represent only a subset of values; and (4) make little use of intuitive judgments and hunches.

To be sure, the potential virtues of rational-analytic thinking are many and well documented, even given the need to be realistic about how well options can be evaluated given the issues of bounded rationality. Further, the more intuitive approaches to decision making may be psychologically attractive, but—to paraphrase—it must be said that when "such an approach is good, it can be very, very good, but when it is bad, it can be horrid." It seems clear to us that decision making and related decision support need to draw on both approaches.

3.2 Working with both rational-analytic and naturalistic styles

With such considerations in mind we have been contemplating ways to achieve a synthesis of theories and styles, one that would avoid the most egregious errors of both, and prove more effective in DSS than the pure rational-analytic paradigm.[16] Here we offer new suggestions, specifically for DSS. Our suggestions are as follows (see also Figure 1):

- Assure that risks and risk mitigation are covered effectively, but provide multiple mechanisms for doing so, so as to deal with different decision maker styles. In essence, we are creating our own FAR strategies for dealing with risk.
- Serve a rational-analytic decision maker with decision support that highlights risks and ways to mitigate them up front, as part of his choice of approaches (a breadth-first approach). In doing so, however, and as a distinct departure from what has often been done, present alternatives that are not "pure strategies" suitable for particular futures, but rather FAR strategies.
- Serve a more intuition-driven decision maker by developing plans seeking to implement "his" COA, but–in the process—identify potential hurdles to be overcome and contingencies against which to prepare hedges. In doing so, employ methods such as assumptions-based planning (ABP)[17]—*not to "fight the decision*," but to identify what might go wrong and prepare accordingly.
- In both cases, deal with all relevant factors, whether hard and objective, or soft and subjective. "Dealing with" may mean creating slots or other opportunities for decision makers to insert their own judgments, rather than attempting to provide staff-generated judgments on matters beyond their expertise, but the factors should appear naturally, rather than being omitted.

The approach we suggest has basis. Empirical research has found that intuitively driven entrepreneurs are, in fact, interested in risks, and also distinguish sharply between risk-taking and gambling.[18] However, their interest is in identifying risks so that they can be overcome, whether by steamrolling over them or adopting various tactics such as buying out or undercutting competitors. It is interesting that while the Bush Administration's approach to the war in

Iraq reflected strong intuitive considerations, there was a great deal of preparation for things that could go wrong. What *did* go wrong had not been adequately prepared for, but the point is that even leaders who had chosen a course of action and did not like broad discussion of alternatives were quite willing to deal with at least some of the risks.

3.3 How much uncertainty analysis can decision makers deal with effectively?

A next issue is how much uncertainty analysis the decision maker can deal with effectively, whether it be through rational-analytic processes or, somewhat more indirectly, through the kind of naturalistic approach allowed for in the right side of Figure 1. A conclusion from past RAND work is that confronting issues of uncertainty well is so difficult that decision makers will do well if merely they consider uncertainty to the extent suggested by Figure 2, which is purely notional. It is expressed in the alternative-COAs format of the rational-analytic decision maker, but the options

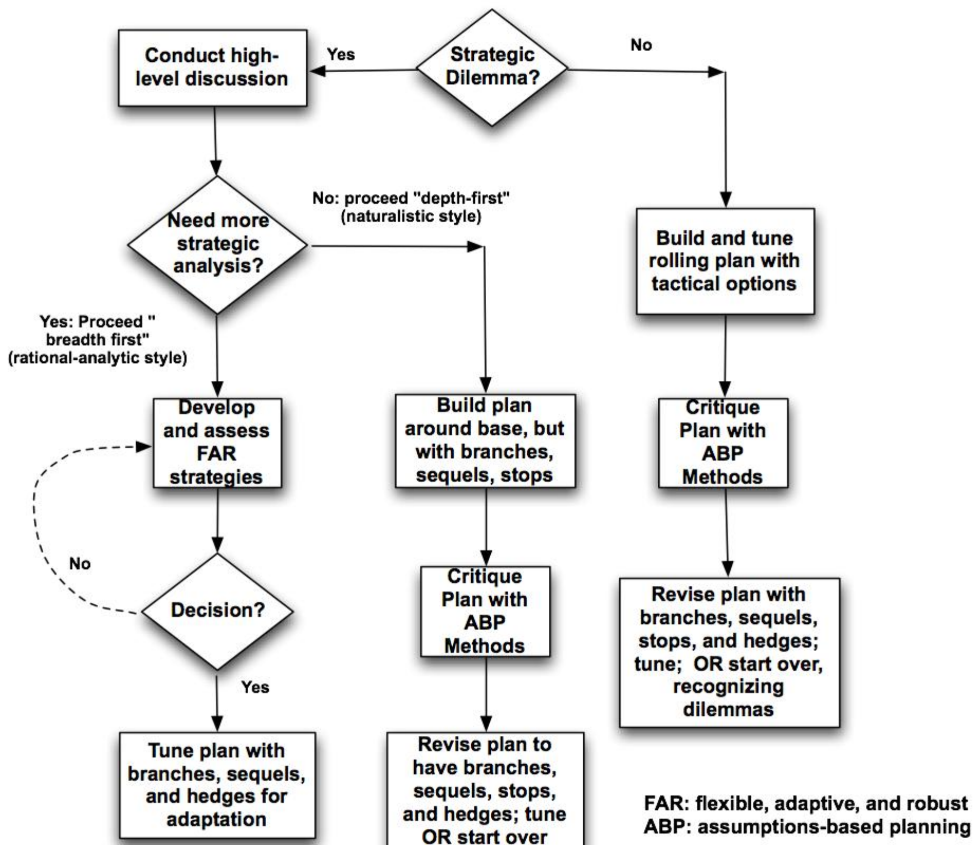

Figure 1: A Flow Model for When To Use Different Approaches in Choosing Options and Building Plans

might be expressed as "The basic plan," "Adjustment 1," and "Adjustment 2" if working with a more intuitive decision maker willing to tolerate adjustments that improve odds of success. The decision support here consists of showing decision makers evaluations of options that estimate most likely, best-case, and worst-case outcomes. The net assessment is left to the decision maker because his own style and various contextual issues outside the scope of the decision-support group will determine whether, for example, he is more or less risk-averse. An important part of this approach is the requirement that the decision maker be able to review the underlying analysis. If, for example, he did not expect the "worst outcome" for COA 3 to be "Very Bad," but rather something more like Marginal or "sort of" Bad, then he should be able to review the logic by zooming to the second part of the figure, which provides the logical rationale—for COA 3—for evaluations of "Worst Outcome" as a function of several factors (analogous tables would apply to the other COAs). As it happens, for COA 3, the analysis reveals logical case 4 as the outcome, for which both Factors 2 and 3 are assessed as Bad and the Worst Outcome is said therefore to be Very Bad. The underlying logic for this assumes that two Bads would lead to a Very Bad, regardless of other considerations. If the decision maker disagreed, he could intervene. The framework for doing so would be relatively clear. Interestingly, this approach of simple decision and logic tables arose from 1980s work by one of us (Davis) building artificial-intelligence models to represent decision makers in simulation. That work stimulated aspects of subsequent work in, e.g., influence-net research. In more recent years, historical research on national crisis decision making has tended to confirm that human decision makers strive for something at roughly the complexity of Figure 2's top portion (unpublished work by Michael Egner and the senior author).

Course of Action	Most Likely Outcome	Best Outcome	Worst Outcome	Net Assessment
COA 1	Good	Good	Good	?
COA 2	Good	Very Good	Marginal	?
COA 3	Good	Very Good	Very Bad	?

Logical Case	Factor 1	Factor 2	Factor 3	Worst Outcome
1	Good	Good	Good	Marginal
2	Good	Good	Bad	Bad
3	Good	Bad	Good	Bad
4	**Good**	**Bad**	**Bad**	**Very Bad**
5	Bad	Good	Good	Bad
6	Bad	Good	Bad	Very Bad
7	Bad	Bad	Good	Very Bad
8	Bad	Bad	Bad	Very Bad

Note: Factors 1, 2, and 3 can have values of Good or Bad; "worst outcome" can have values of Marginal, Bad, and Very Bad. The logical cases are merely the eight different combinations of Good and Bad for the three factors.

Figure 2—A Generic Framework for Decision Making Amid Uncertainty

The display in Figure 2 does not indicate how the various columns are to be weighted because, by and large, we believe that decision makers should see the factors and do their own integration, especially because staff analysis on this crucial matter is often not synchronized with the decision maker's thinking until late in the process. This said, it is easy to insert a "slot" with weights on the three cases, perhaps with perhaps with alternative results for positive-leaning and more conservative decision making attitudes.

Significantly, judgments about probability, which are already folded into the assessments, are often correlated with ambitions (related to "motivated bias") and judgments about the ability to "make things probable" by one's own actions, such as rapid decisive operations. That style of thinking has surely been characteristic of many "great leaders," for better or worse. Napoleon, despite many glorious victories, returned from his winter assault on Russia with only a small fraction of his army still alive. It is not that Napoleon was "wrong" in seeking a glorious outcome, but that he—as many before and after him—overestimated the plausibility of that glorious outcome and did not pay adequate heed to the risk of catastrophic defeat. Those tasked with supporting decisions should bear this in mind and attempt to assist the decision maker in disentangling notions about probability and personal values. As another observation, in this approach the word "utility" does not appear, although it is something that could be inferred after decisions are made. A decision maker who tilts toward weighing the upside, beyond what is justified by estimates of probability, apparently sees a larger magnitude for the positive upside outcome than for the negative downside outcome.

Such weightings relate not to subjective probability, since those are used in reaching the assessments in the first place, but rather to the relative significance placed by the decision maker of gains and losses. That personal weighting scheme is a function of decision maker personalities and context, including history. The contextual role has been studied at length in the Prospect Theory of Daniel Kahneman and Amos Tversky (see review[16] for multiple citations).

Such a decision aid contrasting most-likely, upside, and downside outcomes is useful, but nothing like fool-proof. In particular, people are notorious for having greater confidence than is warranted empirically. Even experts will often say that they are 95% sure of something (or that "it's a slam-dunk"), when their knowledge would better justify a figure of something like 80%, or even much less. With such thoughts in mind, we suggest that decision support should routinely include something like Table 1, which assesses the credibility of estimated confidence levels as a function of process. If the assessments are based only on in-group judgments (to include reach-back from the battlefield), a "one-sigma" notion about best- and worst-case outcomes might be moderately credible, but any claim that the outcomes shown were "two-sigma" values (i.e., less than a 10% chance of any worse outcome) should be given low credibility. If, however, the judgments also reflect serious use of Delphi[19] or related methods, drawing upon an outside community of experts (i.e., drawing on the so-called wisdom of crowds, to cite the current popular work on the subject[20]), then the judgment might have moderate credibility. If, in addition, historical evidence supported the judgment, then credibility might be considered high. As an example, had the credibility of the assessment of no stabilization-and-reconstruction phase been tested in this way before the Iraq war, it would assuredly have come back "low." Many general officers were skeptical and historical work suggested the need for many more forces than were being sent. Although there would still have been legitimate reasons for estimating that the conflict would end quickly, the confidence in that estimate would have been deemed low. Related techniques for debiasing confidence have been discussed elsewhere over the years, as in Chapter 9 of a text by Bazerman.[21]

Table 1. Credibility of Claimed Confidence Level in Assessment of Worst-Case (or Best-Case) Outcome

Claimed Confidence Level **Basis of Judgment**	**Moderate (66%, or "one-sigma")**	**High (95%. or "two-sigma")**
Local staff and limited reach-back to like-minded colleagues	Moderate	Low
Above + self-critical reach-back and analysis to gain wisdom-of-crowd benefits	High	Moderate
Above + use of empirical data from roughly analogous situations	High	High

4. A FRAMEWORK FOR DECISION MAKING AND DECISION SUPPORT

Against this background, let us now describe what we see as the requirements for a decision-support framework. We discuss, in turn: (1) option development and a (2) a portfolio-management approach to option evaluation. Under the latter, we discuss a top-down mechanism of evaluation, treatment of risk within that approach, cost-effectiveness analysis, and the mathematics and logic of aggregation, which is important throughout the portfolio-management approach.

4.1 "Serious" options, not stereotypes

In quantum mechanics and mathematics, defining the orthogonal variables of a problem is powerful and illuminating. However, in the world of strategic decision making, using orthogonal strategies as options can be positively mischievous. It is one thing to use orthogonal dimensions to define a "scenario space" or a space of possible strategies, but the wisest strategies will typically not be "pure," but mixed. There are at least two reasons. First, strategic decision makers often have multiple objectives; thus, the intent should be to find strategies that do reasonably well against all of the objectives. Second, as discussed in earlier sections, the generic solution to dealing with profound uncertainty is flexibility, adaptiveness, and robustness of strategy (see also work by Lempert and colleagues).[22]

The parallel in operations planning might include adopting a strategy that seeks to accomplish multiple goals with two or more tacks being taken so as to improve the probability of success. This might include heavy preparation of the battlespace with air power *and* ample ground-maneuver forces *and* information operations *and* taking every effort to avoid unnecessary collateral damage. The emphasis on these would be uneven, but all would be included.

All of this may seem obvious, but the emphasis on orthogonal strategies is often taught in the classroom, where apparent clarity is sometimes perhaps prized over practicality. Further, we observe that there often exists in real-world decision-making settings a tendency to "tilt" toward a pure solution, something we have observed first hand in our work with government organizations. Resisting such tendencies is important and, we might add, consistent with the instincts of top-flight decision makers.

The first leg of the framework, then, is an attitude of emphasizing alternative ways to achieve FAR strategies from the outset.

4.2 Portfolio management tools

4.2.1 General requirements

The second part of our suggested approach is to use portfolio-management methods and tools to evaluate candidate strategies for their "FARness." Our thinking on this matter stems from considerable work over the last decade.[23-25] In recent work for the Under Secretary of Defense for Acquisition, Technology, and Logistics we have identified key elements of what we see as a generic analytical framework for reviewing capabilities, in and across capability areas, and for making economic tradeoffs among options. These include: (1) routine use of portfolio-management tools; (2) availability of two or more levels of zoom to so that the basis for top-level assessments can be reviewed quickly, and so that a kind of rigor can be maintained; (3) parametric capability models for systems analysis; and (4) an underlying family of models, games, experiments, and historical analyses.

A crucial objective is to provide a "strategic" assessment that is comprehensible at a glance, yet comprehensive in touching upon all the "critical components," a core concept of capabilities-based planning and system engineering. One key feature is candor—showing aspects of both failure and success. Routinely achieving candor is often difficult, especially in environments characterized by obsession with consensus or with protecting favored programs. Success depends on the environment established by leadership. The portfolio management tools should make it easy not only to see gaps, but to help decision makers decide how to adjust the portfolio so as to fill gaps, balance risks and opportunities, prioritize by groups rather than by discrete activities, and even to conduct investment analysis, such as marginal or chunky marginal analysis.

Such ambitions cannot be achieved with a single level of detail, such as a top-level stoplight chart. Decision makers must insist on levels of zoom or "drill-down." Without this, there is little basis for understanding, or for knowing how best to question and adjust assumptions. In practice, decision makers themselves will do only limited zooming, as in spot-checking, but the results will strongly affect their confidence in the quality of the staff work and programs. The expectation of such spot-checking will greatly enhance the rigor of staff work and the clarity with which messages are communicated.

The analytic framework for such assessments should depend not just on one or two models, but on entire families of models, human war games, experiments, and other forms of analysis such as historical studies.[26] Commissioning the assembly of such families, whether centralized or virtual, should be an explicit management action. Success will require drawing heavily on cutting-edge theory and practice in multiresolution modeling and simulation.[27] The idea of families of tools has been discussed extensively in several forums, including a Senior Advisory Group on DoD's modeling and simulation, a recent workshop held to support development of an M&S Master Plan for analysis, and a National-Academy study on DoD's M&S.

4.2.2 Treatment of risk within a portfolio-management DSS

The next question is how to represent risk within a portfolio-management decision support system (DSS). A purely generic taxonomy has eluded us because so many differences exist among application areas. However, as an example, in our current work for OSD (AT&L) we have been considering alternative ways to achieve a particular mission capability requiring acquisition of new systems or new ways of employing existing systems. For that purpose we found ourselves distinguishing among (1) acquisition risks (feasibility, programmatic, and political-stability risks); (2) at-the-time strategic risks (foreign and domestic support); (3) operational risks (including control of effects); and (4) subsequent strategic-effect risks.

If an operational commander, such as a JFACC, were considering options in wartime, a similar taxonomy would apply, but would not include acquisition risks and would include increasingly detailed sublevels of operational risk.

4.2.3 Places to reflect risk in a portfolio management DSS

From a technical perspective, the next issue is where and how to represent the various risks in a portfolio-oriented DSS. The issue is nontrivial because the basic concept is to proceed top down and achieve comprehensibility. As one thinks of more important matters to represent, the effect can easily be to overwhelm the decision maker with more and more complex tables and graphics. Furthermore, it is inappropriate to represent all of the various risks in the same way because they vary in their nature, significance, and analytical status. Although this is also an ongoing activity of discovery, the methods include: (1) measuring effectivenss in "bad cases" as well as nominal cases; (2) using the graphic equivalent of cautionary footnotes; (3) including explicit measures of composite risk; (4) peforming underlying calculations of effectiveness with safe-sided assumptions.

4.2.4 Lowest-level explanations

The next technical issue that we have grappled with is how to provide immediate explanations of analytical results summarized in high-level decision support when the questions asked become even more detailed and "technical." Although we have referred above to zooms, the nature of what is possible and appropriate depends on how deep one wishes to go and what the nature of the issue is. A full discussion of this matter goes far beyond the scope of the present paper, but it is clear that to the maximum extent possible, essential documentation of both assumptions *and logic* should be included within the DSS and should be understandable to the decision maker's top analysts. This can *usually* be accomplished adequately with a combination of intuitively understandable variable names, variables, *well-structured* assumption lists, some overview graphics such as "live" exploratory-analysis charts allowing interactive response to questions, and simple logic tables.

4.2.5 The mathematics and logic of aggregation

One of the most difficult technical issues in designing top-down decision support is that high-level displays are abstractions of lower-level knowledge. What abstraction methods should be used, and when does it matter? The first observation here is that, even though linear-weighted sums are often very good aggregation mechanisms, they can also be quite misleading. For example, if a system's capability is being assessed and that system depends on each of several critical components performing adequately, then a linear approximation would mistakenly suggest that one could compensate for failure of one component by buying more of another.[14] We are currently using a variety of non-linear aggregation methods involving thresholds, critical-component effects, probability-related mathematics, and S-curves.

4.2.6 Cost-benefit information and chunky marginal analysis

The last item in our requirements for features in a portfolio-management tool relates to cost-benefit calculations and the related issues of marginal and chunky-marginal analysis. Since any composite measure of overall effectiveness will be highly suspect because of the uncertainties or disagreements about how to aggregate across different objectives and measures, it is essential to have a mechanism for exploring the consequences of different "perspectives" about, e.g., the relative importance of different missions and constraints, the relative probabilities of various risks coming in to play, and so on. For good reason, single-assumption-set analysis of effectiveness and cost effectiveness will likely have little influence on savvy decision makers.

There is need for both marginal analysis, identifying where to spend or cut the marginal dollar (or billion dollars), and a more chunky type of analysis that uses larger increments of spending (or cuts) and is able to account for S-Shaped phenomena such as where a high-payoff system requires a huge up-front investment and time before any payoff is achieved.

Finally, there is need for cost-benefit comparisons on large composite options, such as alternative defense programs or alternative POMs (the yearly Service program submissions). This is not marginal analysis, but a more strategic level of cost-benefit work. The classes should be related, however, in that the composite options examined seriously should be motivated by more microscopic analysis; they should be the best-of-breed options for alternative strategies.

5. EXAMPLES USING NEW TOOLS

In this section we illustrate some tools that we see as quite consistent with the spirit of this discussion. These relate to portfolio analysis, to capabilities-based planning, and to systems analysis.

5.1 Portfolio analysis

The Portfolio Analysis Tool (PAT) is a generic offspring of a tool (PAT-MD) developed earlier for the specific purposes of the Missile Defense Agency (MDA).[25] It was motivated by earlier work with a similar tool called DynaRank,[24] which has also evolved over time, in part as the result of military, environmental, and health applications, and in part from lessons drawn from work with PAT and PAT-MD. We shall focus here on PAT.

PAT is a spreadsheet application, built in Microsoft EXCEL®. In practice, of course, such a tool stems from a great deal of thinking, programming, experience, and iteration. As with much software, usefulness depends upon details. In this paper, however, we shall focus primarily on major functionalities.

Figure 3 shows a top-level summary chart from a hypothetical PAT application to operations planning. It is assessing three options against a number of criteria). Figure 3 is just a familiar score card, although it appears gray-scaled here rather than being shown in its normal color-coded fashion. Note, however, that it includes a risk column. In this particular case, no "net assessment" is shown, although can easily be added. The context is that a JFACC is considering three courses of action before recommending an approach to the Joint Forces Commander. The first is the JTF commander's clearly preferred option because of a strong sense—based on long experience and intuition—that it is essential that the operation be as coalitional as possible so as to build cooperation, trust, and shared responsibility. That base plan is to do "everything" with the coalition partners involved as much as possible in all activities. The JFACC, however, is concerned that the base plan involves substantial risks. First, the plausibility of attacking the enemy's air forces and air defenses with full success may well depend on achieving surprise, which will be much more complicated and dicey in a fully coalitional attack. Second, the JFACC is not confident of the prowess of all of the coalitional members' operators or of the reliability of their systems. He has no such concerns about some of the allies with whom he has worked previously, but is very wary of assigning important functions to others. The second and third options, then, are mere "variants" of the base plan: there is no attempt to fight the JTF Commander's guiding principle, but rather an attempt to mitigate the risks of the base plan. Option 2 would do that by having the *initial* strike be conducted entirely by US stealthy aircraft and missiles so as to achieve maximum surprise. There would be no "local" indicators that a strike was underway and coalitional forces generally would be ready for battle, but not visibly prepared for imminent battle. Option 2 is deemed likely to disarm the enemy's air forces and air defense systems, making it possible for air forces to provide superb cover for subsequent ground-force maneuver. Option 3 is another variant, this one emphasizing use of only missiles and information-warfare assets during the initial strike. Proponents of that variant emphasize that it would minimize the possibility of coalitional losses and "should" be able to disarm the enemy's air forces and air defenses. The idea in both Options 2 and 3 is that there would be extensive consultation and coordination with coalition members, consistent with the JTF commander's intent, *except* that there would be no prenotification of precisely when the initial strike will occur so as to maximize surprise. Proponents of the options believe that it is likely that the allies would be assuaged by the consultation and coordination, and that the operation as a whole would be fully coalitional.

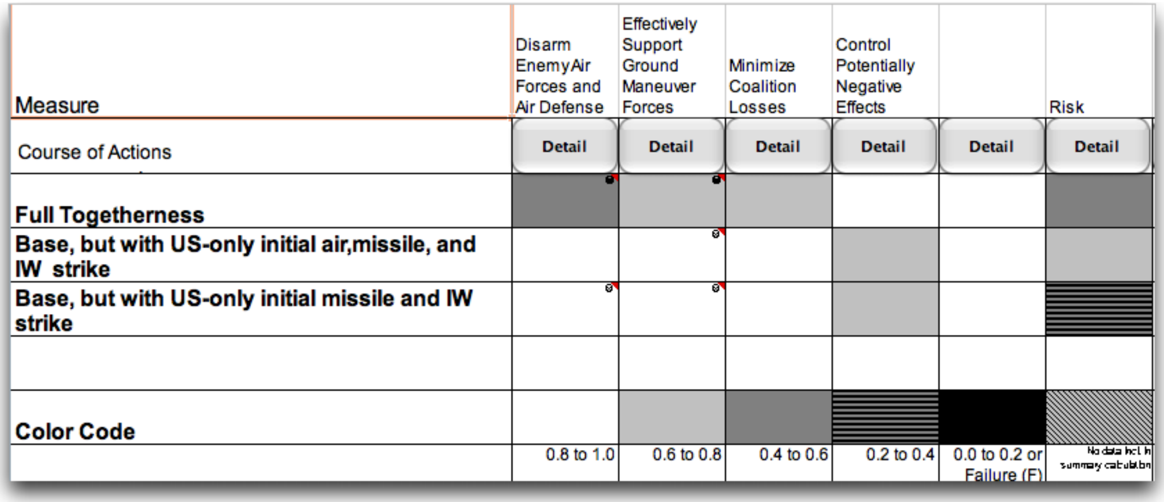

Measure	Disarm Enemy Air Forces and Air Defense	Effectively Support Ground Maneuver Forces	Minimize Coalition Losses	Control Potentially Negative Effects		Risk
Course of Actions	Detail	Detail	Detail	Detail	Detail	Detail
Full Togetherness						
Base, but with US-only initial air,missile, and IW strike						
Base, but with US-only initial missile and IW strike						
Color Code						
	0.8 to 1.0	0.6 to 0.8	0.4 to 0.6	0.2 to 0.4	0.0 to 0.2 or Failure (F)	No data incl. in summary calculation

Figure 3—A Summary Portfolio Display for a JFACC

This particular display does not show an "upside" to balance the "downside" (risk) because the entire operation is already being planned for the upside case of decisive victory. The arrangement in Figure 3 is to show four measures of goodness on the left side (prospects for disarming the enemy's air forces and air defense; supporting ground-maneuver forces; minimizing coalitional losses; and controlling potentially negative effects, such as fracturing of the coalition due to non-cooperation and negative strategic blow-back that might occur if collateral damage were high or the world did not believe that the attack was appropriate in nature. By these standards, the basic Full-Togetherness plan fares rather well in the first and fourth measures (white is very good, by analogy with green in a stoplight chart). It does a bit less well in the two middle categories. The variants with a surprise US-only or US-dominated initial strike, do very well *except* that some tension with alliance members could be anticipated, which might have some temporary repercussions, such as delays in ground-force maneuver or noncooperation of various types. Such problems, however, are assessed to be not very serious (light gray is "good", although not "very good"). Upon looking at the Risk column, however, we see more stark differences. All of the courses of action have risks, but those of the base plan and the second variant are significantly more risky. An "eyeballing" of the chart suggests that the middle option might well be preferred.

By clicking on the "Details" button of any column, one can bring up a zoom, such as a zoom on what constitutes Risk. The result would also be a spreadsheet view, but might show Risks as having been judged from operational-risk and strategic-risk components. Operational risk might be deemed quite high for the baseline option, but low for the first variant (option 2). In contrast, Option 2 might have some strategic risk that the baseline option does not. Why? Because, despite the assessments of the JFACC's staff and various experts they have consulted, the possibility exists that important coalition members will be much more angered than expected when a surprise initial strike occurs about which they were not notified in advance. Although the groundwork would have been layed in advance, so that only the timing would be a surprise, perceptions on such matters are not entirely controllable. Even if military-to-military communications are good, reactions of capitals might be surprising. Thus, there are tradeoffs. The third option, depending strictly on missiles and information warfare, is deemed to be even more inherently risky. Its operational success would depend on the quality of detailed technical intelligence and perceptions of its appropriateness by the world are in doubt because there would be temporary effects that might not be recognized as temporary.

A second level of zoom might be to understand better how operational risks are assessed. Yet again, the result might be a spreadsheet view, but in this case it might distinguish between, say, the assessment of ability to get to the target and ability to disarm the target in the event that it is reached. In actual applications, this level of detail could be much richer.

Although this example has been high simplified, we hope that it conveys the sense of the tool. Actual applications involve more variables, more subtleties, and a good deal of work.

5.2 "Zooming" into underlying systems analysis

The zooms discussed above work rather well so long as the issues can be understood in essentially logical terms, as depicted by a table with variables having discrete values, such as low and high, good or bad. A different type of zoom, however, is needed to understand issues in more detail. One natural format for such detailed material is the kind of parametric chart of which systems analysts and operations researchers are fond. Such charts, with overlays, can explain the judgments reflected at higher levels with discrete cases, and do so in a broader context that reveals, e.g., phenomena such as sharp boundaries versus asymptotic tails. We do not show this type of zoom here, but we use it routinely with PAT, sometimes using analytical models embedded in PAT itself and sometimes by pulling up another model running at the same time as PAT. We also do not show here displays relevant to resource allocation, but for comparing investment options it is easy to construct cost-effectiveness plots that visually indicate where the payoff is highest for different marginal or chunky-marginal investments.

5.3 Marginal analysis

Although we shall not elaborate here, PAT has been used primarily for investment analysis to date. In those applications, the options are not operational courses of action, but alternative investment programs. An illustrative output from such a study is a scatter plot of effectiveness versus expenditure, with each point being a different option. It is quite easy to see from such a plot where to spend the marginal dollar (or billions of dollars)—if the composite measure of effectiveness is satisfactory. Since the measure is very likely controversial, we explore the significance of different perspectives of how to calculate effectiveness. Often, some options are robustly preferable because they have significant albeit uncertain value and don't cost very much—although they may raise organizational hackles or require changes in the way business is conducted. In other cases, relative cost effectiveness does indeed depend on perspective and analysis cannot resolve the disagreement. Policy makers must, for example, balance near-, mid-, and long-term considerations. They must also assess subjectively the relative credibility of options, based, for example, on their advocates' record of performance.

6. CONCLUSIONS AND NEXT STEPS

In this paper we have sought to describe the rationale behind an approach to high-level decision support. We have also illustrated how it can be implemented, using recently developed RAND tools. As always, however, our successes have also raised new challenges for research and the development of enabling technologies. We mention the following in particular: (1) Developing new models and tools designed specifically to assist in finding flexible, adaptive, and robust strategies; (2) Creating subtools making it much easier to develop and document alternative aggregation and disaggregation rules on the fly, tools that would allow thinking to be done mathematically, graphically, or logically; (3) Creating tools for "exploratory analysis" at the portfolio-structure level, i.e., tools allowing the analyst to generate an appropriately wide range of "perspectives," compute the corresponding displays, and abstract from this exploration a compact and insightful summary of what has been learned; (4) Creating network-based tools to facilitate reach-out exercises, such as those using Delphi or other methods for tapping into group wisdom in a sound manner. This would involve mechanisms for communication; expressing problems; identifying appropriate samples of experts (and perhaps others); collecting the results; aggregating in ways that preserve information about contrary views and usually ignored possibilities; and reflecting the results in high-level decision support; (5) Developing techniques for readily creating and consistency maintainence among multiresolution structures for decision-support hierarchical displays; and (6) Developing techniques and protocols making it easier to connect portfolio-style tools to models of differing character and resolution, rather than having to rely upon both sneaker-net methods and extensive side analysis.

ACKNOWLEDGEMENTS

This paper builds on research sponsored by the Office of the Secretary of Defense and the Air Force Research Laboratory (Rome, New York). It also reflects thinking done by one of us (Davis) as part of a National Academy of Sciences study on DoD's modeling, simulation, and analysis. We wish to acknowledge the collaboration in the research discussed here by colleagues James Bonomo, Paul Dreyer, Richard Hillestad, Russell Shaver, and Henry Willis.

REFERENCES

1. Kent, Glenn A., and William E. Simons, "Objective-Based Planning," In *New Challenges in Defense Planning: Rethinking How Much is Enough*, edited by Paul K. Davis, Santa Monica, Calif.: RAND Corporation, 1994.
2. Deptula, Brigadier General David, "Effects-Based Operations: Change in the Nature of Warfare," http://www.afa.org/media/reports/, (as of 2001).
3. Davis, Paul K., Effects-Based Operations: A Grand Challenge for the Analytical Community, Santa Monica, Calif.: RAND Corporation, 2001.
4. Dubik, James, "Effects-Based Decisions and Actions," *Military Review*, Vol. January-February, No., 33-36, 2003.
5. Smith, Edward A., *Complexity, Networking, and Effects-Based Operations,* CCRP, Washington, D.C, 2006.
6. Rayburn, MG Bentley (USAF), "Building an Assessment Structure, briefing at Conference on Effects-Based Operations, Tysons Corner, VA, Vol. January 31-February 1," 2006.
7. Jobbagy, Zoltan, "Theory, Reality and the Nature of War: Effects-Based Operations Meet Post-Conflict Operations", Clingendael: Centre for Strategic Studies, 2006.
8. Joint Warfighting Center, Commander's Handbook for an Effects-Based Approach to Joint Operations", Suffolk, VA: U.S. Joint Forces Command, 2006.
9. U.S. Air Force Air Materiel Command, Commander's Predictve Environment: Understanding the Battlespace, 2006.
10. Kahan, James P., D. Robert Worley, and Cathleen Stasz, *Understanding Commanders' Information Needs,* Santa Monica, Calif.: RAND Corporation, 2000.
11. Botterman, M. et al., Foresight Cyber Trust and Crime Prevention Project: Gaining Insight From Three Different Futures, London: Office of Science and Technology, 2004.
12. Davis, Paul K., "Planning for Adaptiveness," In *New Challenges in Defense Planning: Rethinking How Much is Enough*, edited by Paul K. Davis, Santa Monica, Calif.: RAND Corporation, 1994.
13. Davis, Paul K., David Gompert, and Richard Kugler, "Adaptiveness in National Defense: The Basis of a New Framework," Santa Monica, Calif.: RAND Corporation, Issue Paper IP-155, 1996.
14. Davis, Paul K., Analytic Architecture for Capabilities-Based Planning, Mission-System Analysis, and Transformation, Santa Monica, Calif.: RAND Corporation, 2002.
15. Light, Paul C., The Four Pillars of High Performance, McGraw-Hill, 2004.
16. Davis, Paul K., Michael Egner, and Jonathan Kulick, Implications of Modern Decision Science for Military Decision Support Systems, Santa Monica, Calif.: RAND Corporation, 2005.
17. Dewar, James, Assumptions Based Planning, Cambridge, London, UK, 2003.
18. March, James G., with the assistance of Chip Heath, A Primer on Decision Making: How Decisions Happen, New York, NY: The Free Press, 1994.
19. Linstone, Harold A., and Murray Turoff, eds, *The Delphi Method: Techniques and Applications,* Newark, NJ: New Jersey Institute of Technology (available on-line), 2002.
20. Surowiecki, James, *The Wisdom of Crowds,* Anchor, 2005.
21. Bazerman, Max H., *Managerial Decision Making (2nd ed.),* John Wiley, New York, 1986.
22. Lempert, Robert J., Steven W. Popper, and Steven C. Bankes, *Shaping the Next One Hundred Years: New Methods for Quantitative Long-Term Policy Analysis,* Santa Monica, Calif.: RAND Corporation 2003.
23. Hillestad, Richard, and Paul K Davis, *Resource Allocation for the New Defense Strategy: The DynaRank Decision Support System,* Santa Monica, Calif.: RAND Corporation, MR-1996-OSD, 1998.
24. Dreyer, Paul, and Paul K. Davis, *The Portfolio Analysis Tool for Missile Defense (Pat-Md): Methodology and User Manual,* Santa Monica, Calif.: RAND Corporation, TR-262-MDA, 2005.
25. Davis, Paul K., "New Paradigms and New Challenges, keynote presentation, Military Track," *Proceedings of the Winter Simulation Conference*, 2005.
26. Davis, Paul K., James H. Bigelow, and Jimmie McEver, *Exploratory Analysis and a Case History of Multiresolution, Multiperspective Modeling,* Santa Monica, Calif.: RAND Corporation, RP-925, 2001.