REPRINT

Modeling Adversaries and Related Cognitive Biases

JONATHAN KULICK, PAUL K. DAVIS

Reprinted from *Proceedings of SPIE*

RAND
PROJECT AIR FORCE and
NATIONAL DEFENSE RESEARCH INSTITUTE

The RAND Corporation is a nonprofit research organization providing objective analysis and effective solutions that address the challenges facing the public and private sectors around the world. RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

Published 2003 by the RAND Corporation

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

Preface

This reprint volume of work in progress collects three papers related to cognitive modeling of adversaries. The first, "Synthetic Cognitive Models of Adversaries for Effects-Based Planning," describes a top-down theory-driven approach, its motivation, and past examples. One of its themes is the importance of having alternative models to force recognition of uncertainty. The second, "Thoughts on Higher-Level Adversary Modeling," extends the discussion and specifically addresses the need, in high-level decision support for effects-based planning, to keep such adversary models extremely simple and to use them to improve assessment of best-estimate, best-case, and worst-case outcomes for alternative courses of action. The third paper, "Judgmental Biases in Decision Support for Strike Operations," is a broad discussion of judgmental biases in decision support and efforts ("debiasing") to mitigate those biases; the paper includes discussion of biases that affect a side's mental image or model of the adversary, as well as giving speculative examples relevant to planning of air operations. The first paper was fully sponsored by the United States Air Force Research Laboratory; the second and third also benefited from research for the Defense Advanced Research Projects Agency.

Synthetic Cognitive Modeling of Adversaries for Effects-Based Planning

Paul K. Davis

Reprint from proceedings of SPIE's AeroSense 2002 Conference

Reprinted with permission from *Proceedings of SPIE: Enabling Technologies for Simulation Science VI*, Vol. 4716, April 2002, edited by Alex Sisti and Dawn Trevisani, "Synthetic Cognitive Modeling of Adversaries for Effects-Based Planning, Paul K. Davis, pp *236-250*.

Synthetic Cognitive Modeling of Adversaries for Effects-Based Planning

Paul K. Davis
RAND, 1700 Main Street, Santa Monica, CA 90407-2138

ABSTRACT

Models of adversaries' reasoning can be constructed to inform development of adaptive strategies, including strategies that include effects-based operations. Such models can apply to individual leaders or to groups that one seeks to influence. This paper describes an approach to building such models. The results are top-down, highly structured, driven by theory, designed with multiresolution methods that permit "zooming in,, on issues, and suitable for use in high-level decision meetings. The models have been qualitative and non-automated, but the methodology could usefully be incorporated into a more general computer-supported decision-support environment, where it would supplement other tools for decision support.

Keywords: Cognitive models; synthetic cognitive models (SCMs); qualitative models; multiresolution modeling (MRM); model abstraction; multiresolution, multiperspective modeling (MRMPM); decision models; effects-based planning; effects-based operations (EBO); behavioral modeling; decision support

1. INTRODUCTION

1.1 Background

Effects-based operations (EBO) are defined here as operations conceived and planned in a systems framework that considers the full range of direct, indirect, and cascading effects—effects which may, with different degrees of probability, be achieved by the application of military, diplomatic, psychological, and economic instruments.[1] Often, the purpose in effects-based operations is to influence the *behavior* of adversaries—either individuals, such as a Saddam Hussein, or groups, such as the Al Qaeda terrorist group. The influence sought on individuals is often described in terms such as to dissuade, deter, or compel. Actions may similarly influence whole groups, but they may instead be described in terms such as demoralization of an army or causing the collapse of support for the adversary's government.

Although effects-based planning (EBP) was initially regarded by many observers as yet another fad, it is increasingly recognized as central to sound development of strategy. Moreover, it is no longer an abstraction. Not only did the United States practice EBO in the war over Kosovo, so also did Osama Bin Laden and his Al Qaeda terrorist organization practice EBO against the United States when they launched a grandiose attack on symbolic pillars of U.S. capitalism and military strength, killing thousands of innocents by intention. Historians will argue for years about the effects on Slobodan Milosevic of the pointed U.S. bombing that characterized the later weeks of the Kosovo war, but no one can doubt the effects on U.S. behavior subsequent to September 11. Those effects were in part intended by Bin Laden (direct killing and destruction, followed by America's shock as it recognized that it was not a sanctuary), in part hoped-for (consternation and enormous efforts to improve security), and in part a very unpleasant surprise (the massive and relentless U.S. war against Al Qaeda and the Taliban in Afghanistan).

As this was written in spring, 2002, world attention focused on the Middle East, where Palestinians had resorted to effects-based operations involving suicide bombers against which Israel had no good defense. Targeting was not against material things, but the minds of Israelis and the world. Would this new strategy succeed and lead everyone to conclude that terrorism works? Or would it eventually fail and lead everyone to conclude that it is ultimately destructive to everyone? What options would Israel have in responding? And what would be the implications for the U.S. war on terrorism? History was being written and the outcome was unclear. It did seem evident, however, that nothing is straightforward in EBO and that developing effects-based strategies is by no means a well-understood military art or science.

Enabling Technologies for Simulation Science VI, Alex F. Sisti, Dawn A. Trevisani, Editors,
Proceedings of SPIE Vol. 4716 (2002) © 2002 SPIE · 0277-786X/02/$15.00

1.2 Synthetic Cognitive Modeling of the Adversary

Against this background, it is plain that a key element in any such art and science of EBO should be influence modeling, so that speculations about influences can be turned into analytical constructs susceptible to critical review and evolution. Some research on influence modeling is now under way—including for effects-based operations. Some research uses techniques such as dynamic Bayesian networks and influence nets, perhaps with significant simplifying approximations, game theory, or a combination.[2,3] This paper describes a different approach, called synthetic cognitive modeling (SCM).

SCM has interesting overlaps with the other approaches, but it is significantly different in both theory and practice. Figure 1 suggests some of the differences. Model A denotes an "empirical„ model to predict behavior as an interpolation or extrapolation from historical events (e.g., the strategic bombing of Germany in World War II, of Iraq in 1991, or of Serbia in the late 1990). Even Model A is a major improvement over unstructured expert opinion. Model A might be constructed from statistical analysis of historical events characterized by enough situational factors to make the model meaningful. Model B would be similar in function, but would be based on "structural theory„ as constructed from history and international relations theory. Such a model would address the personality of the target only vaguely, if at all (e.g., a situational factor might be that the government of the targeted nation was a dictatorship). It would focus on the factors that so-called structural realists worry about, plus some from idealist or organizational-behavior schools (e.g., the subject's social cohesion), but would include nothing psychological or personal. Indeed, Model B would not even be culture-specific. Model C would reflect the approach of a mainstream decision scientist. The model might be conceived in terms of a Bayesian net with conditional probabilities, or an approximation. The diagrams indicating postulated influences might be quite complex with tens or scores of nodes and arcs. The output would be a putative probability of the desired behavior, as computed by an underlying inference engine or a more approximate mechanism.[2,3] The result might or might not have face-level reasonableness, but would likely not be accompanied by much of a substantive explanation. In the extreme, results would have to be interpreted by the infamous expression, "Well, there's a lot going on, but this is what the model says.„ The analyst might be able to do much better than this, by pointing out which influences were triggered by the circumstances, but it would be difficult for an observer to judge the result well because neither the model nor the analysis would be structured for such purposes. What factors have been omitted? What assumptions have been buried? In practice, Model C would also be likely to have been developed for a so-called best-estimate version of the adversary modulated by the use of probabilities ("Well, if ...happened, he'd be 70% likely to ...„).

Finally, Model D suggests the approach of synthetic cognitive models (SCMs), which would attempt to represent alternative reasoning patterns of the adversary. However, it would be attempting to represent the *effect* of that reasoning, not the dynamics. Real humans, after all, ingest and process information in a bewilderingly complex and unstructured way. Only at the time of decision may structure emerge, as when the decisionmaker summarizes for his advisors by saying,

> "Well, as I understand it, we really have only three options. The first option fails because...; the second option fails because...; but the third option stands a chance. We'll go with that.„ The phrases following the "becauses„ would encapsulate key factors in the decisionmaker's reasoning.

As a whole, this summary might have the character of a structured, logical, and rational approach. It would be "as though„ the decisionmaker's reasoning had been that way all along. It is this net assessment that the synthetic cognitive models attempt to capture. The plural applies because a central feature of the approach is to treat uncertainty about the adversary explicitly, with discrete variants of the model. As Figure 1 indicates, Model D is highly structured and theory-driven. Further, it is an example of multiresolution, multiperspective modeling (MRMPM).[4] Thus, if one wants to understand the predictions of a given adversary model, one could get a top-level summary (along the lines of the decisionmaker summary imagined above) or one could go into more details, asking "But *why* is Option 2 too risky?„ Such explanation might go through multiple levels of detail.

Any of the approaches could be modified to take on more of the attributes of its competitors. All of the approaches are attractive. For example, that of Model A has an important place because, as statisticians and historians are apt to point out, reality is often not what the theorists claim. Sometimes, we need to have our noses rubbed in the data. The diversity of approaches is also useful because needs vary with the purpose of the model and context (e.g., decision support for an airwar commander or for the President). In any case, the remainder of this paper sketches the elements of synthetic cognitive models.

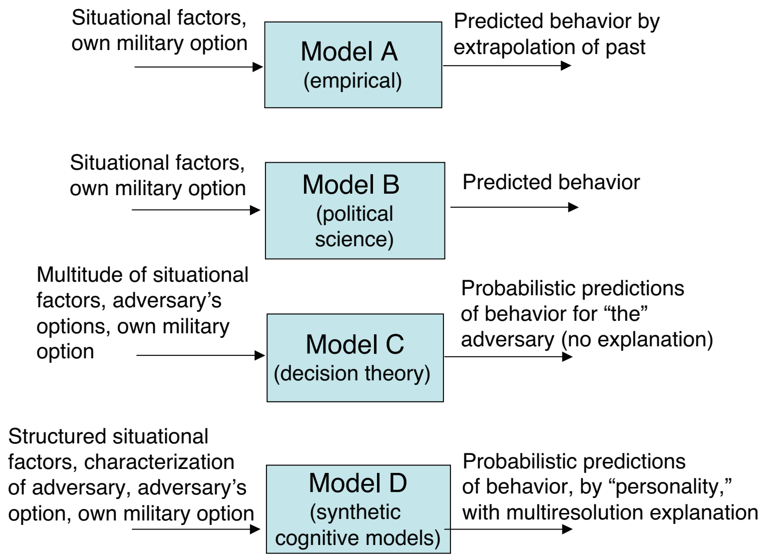

Figure 1—Alternative Modeling Approaches

2. PRIOR WORK ON SYNTHETIC COGNITIVE MODELS

Reviewing the history of the proposed approach is worthwhile because that history contains lessons for the current era, especially for those interested in decision support systems.

2.1 The RAND Strategy Assessment System (RSAS)

The roots of the approach are in work during the 1980s that led to the RAND Strategy Assessment System (RSAS), a very large and complex (million lines of code) system for military and political-military war gaming.[5] The RSAS could be run automated or with decision models (agents) at one or more position. It included Red, Blue, and Green agents representing, respectively, the Soviet Union, the United States, and third countries. Functionally, these were artificial-intelligence (AI) models, although unusual ones.[a] The Red and Blue agents "worried,, about cosmic matters such as whether and when to go to war, escalate, or to terminate. They also worried about what military strategies to adopt and when to change strategy.

The RSAS development was an intellectual and technical success as of 1988, but the Cold War came to an end and support for the effort quickly diminished. The combat-modeling aspects of the system still exist (as the Joint Integrated Contingency Model, JICM), but the agents slipped into history. Much had been learned from their development, however, and that has affected research and analysis to this day.

One lesson related to the RSAS' inherent complexity. Despite *extraordinary* efforts to design it for comprehensibility, it was overwhelming to most would-be users. The system as a whole was almost never used in its entirety outside of RAND, and by only a few of us within RAND.[b] Fortunately, we had anticipated this and designed the system to be

---

[a] Usual AI models were expert systems with a large body of production rules processed by an inference engine written in LISP or PROLOG. After initial experiments with such models, we decided instead to use highly structured top-down models written in a high-level procedural language developed by Edward Hall and Norman Shapiro.[6] The benefits included explanation capability and orders of magnitude in performance. Disadvantages included brittleness—a casualty of not then being able to express the relevant knowledge in declarative terms.

[b] This was discussed in a Yale University dissertation on war gaming and the RSAS experience by John Thomas Hanley (circa 1990).

modular. Thus, most RAND studies and all outside users (e.g., the war colleges) used only the combat- and strategic-mobility simulation, or a subset suitable to a particular theater. Human teams were typically preferred over agents, although they sometimes used and adapted off-the-shelf war plans developed for the military-level agents. Some studies within RAND used one or more of the political-level agents, with only a minimalist stub and interface representing military operations.

Despite its nominal allure, then, full-system operations were rare and not particularly useful. Any notion that the *full* system could be truly maintained and used within the government dissipated. The reasons were many, but my own view was that the price paid to computerize, use, and maintain such ambitious agent-based models was extremely high relative to their value in generating insights. Further, such complexity interfered with clear thinking of the sort so much valued in studies and direct decision support. It might tell us a great deal about reality, but simplification was ultimately necessary. Surely, cream could be skimmed.

2.2 Modeling Saddam Hussein in Strategic-Level Analysis

By the late 1980s, my colleagues and I had turned to new issues, including regional contingencies, but with our RSAS-agent experience behind us.[7] In 1990, before the invasion of Kuwait, John Arquilla and I began developing synthetic cognitive models of Saddam Hussein—being careful to use an approach that could be presented and discussed with policymakers, using nothing more complex than viewgraphs. Our point of departure was the desire for methods encouraging government officials to "play through„ potentially complex contingencies before they arose. In July, 1990, we held a contingency planning war game in which interagency participants were "forced„ to address dilemmas that might arise if Saddam ever threatened Kuwait. After two hours of venting repeating conventional wisdom about how such a threat would be just bluster because Saddam would never invade, participants worked through the what ifs, which were based on an early alternative model of Saddam. A few weeks later, the scenario played out in reality and, on the day of the invasion, a colleague and I were able to quickly write and send to high officials a short paper that laid out the game board and identified dilemmas and options. The value of the prior work, then, was that it had built intellectual capital that could be applied quickly. One participant in the game was a senior intelligence officer, who—influenced by the game—was perhaps the earliest in crisis to conclude that Saddam was in fact about to invade.[Ref.8,pg.217]

Following the invasion, we began in earnest to develop models of Saddam Hussein for a variety of predictable decision points in the months ahead. We used these to write further papers for officials during the war. Subsequently, we documented the work and discussed in some detail why we believed Saddam invaded, refused to be compelled to back out, and then quit suddenly after the U.S. ground counteroffensive.[9,10] Our models of Saddam were later embodied in a computerized decision-support system for war gaming in Europe by the University of the Bundeswehr.

2.3 Subsequent Work in the 1990s

In the process of this research, we had developed a rather generic methodology for modeling adversaries—one based on a combination of strategic thinking, regional studies, gaming, behavioral psychology, and military analysis.[10,11] We (mostly my colleague John Arquilla) later applied the ideas to historical cases with good results.[12] In 1994 we conducted a study on the thinking of potential "proliferators„ and, as part of that effort, held a conference to which we invited academic behavioral psychologists and people who had profiled foreign leaders for the CIA. One purpose of the conference was to expose our approach to this audience to see if we had properly reflected insights from behavioral research. The results were satisfying.[13]

A year later, a colleague and I applied the same methodology quickly to the vexing problem of how to think about North Korea with a new ruler (Kim Jong Il) as the United States developed negotiating objectives. Our analysis was regarded as interesting and provocative, but not, regrettably, persuasive enough.

As a final item of background, a British group led by Paul Willis of the Defense Establishment Research Agency (DERA) adopted and extended some of our methods (particularly those using cognitive maps or influence diagrams) to advise the British government during the late-1990s crisis in over Kosovo. The issue was how to influence Slobodan Milosovic. DERA's work was reportedly quite useful to government officials.

In summary, the ideas sketched in what follows have been applied to a number of problems over the course of a decade with a fair degree of success. The experience has reinforced the view that the top-down, structured, qualitative approach is comprehensible by and stimulating to senior people and their staffs. Let me now sketch key elements of the approach.

3. AN OVERVIEW OF SYNTHETIC COGNITIVE MODELS

3.1 Taking a "Game" Perspective

Basic Structure. The remainder of the paper will focus narrowly on modeling the adversary, but Figure 2 indicates that the context for thinking about the adversary is one in which both sides perceive aspects of the real world and make decisions accordingly. Their perceptions may, of course, be inaccurate. One important aspect of this is that both Red and Blue have models of each other. This can be recursive so that one may have Red, Red's Blue, Red's Blue's Red, and so on. The level of detail needed diminishes rapidly with "depth„ in such a recursion, but history and experience tells us that the third level of depth is significant. For example, General Schwartzkopf understood that Iraq's military believed that U.S. forces would not seriously consider a thrust through the western desert. Red's model of Blue was wrong (due to Iraq's mirror imaging and failure to appreciate U.S. technology, and to U.S. operational secrecy), but Blue's model of Red's model of Blue was correct and important.

Figure 2 may seem unexceptionable, but much work on effects-based operations and rapid decisive operations has been hampered by the artificial separation of work on intelligence and operations within the standard military planning process. It has not been straightforward to organize for thinking about such "game„ issues.

Multiresolution Modeling. It is only natural to think about the adversary hierarchically. Figure 3 suggests what might happen within Red's "Decide Actions module„ if Red were a terrorist group. Even at this simplified level of description, the diagram has a mix of generic elements and context-specific elements. It refers to Red's beliefs about Blue, its assessment of target vulnerabilities, and Red's own capabilities. This assessment, in turn, depends on lower-level (higher-resolution) factors such as the availability of logistics and different mechanisms. This, then, is multiresolution modeling. It is also multiresolution, multiperspective modeling (MRMPM), because anyone using such a diagrammatic approach should be ready to quickly change some of the nodes to highlight considerations from a different point of view. Multiresolution modeling is an essential part of SCM. It is also not as straightforward as it might at first appear, but related principles for accomplishing it are now emerging.[4]

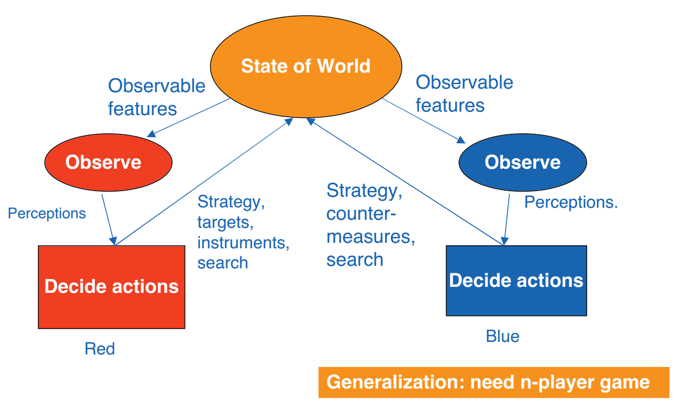

Figure 2—A Simplified Game-Structured View

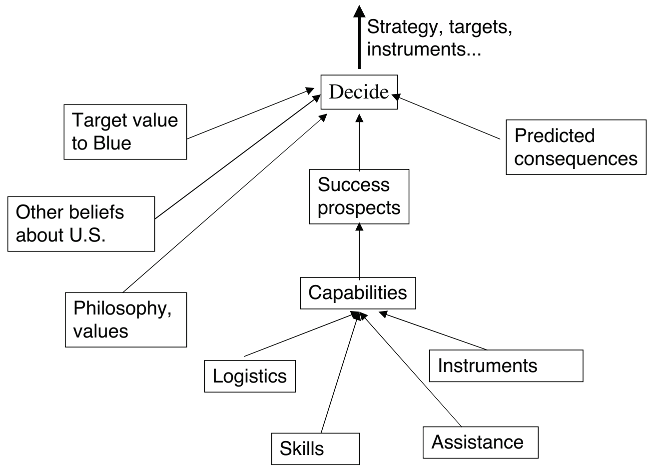

Figure 3—An Illustrative Decomposition of Red's Reasoning (Terrorist Example)

3.2 The Assumption of Limited Rationality

A basic assumption of the approach is that the subject of the modeling exhibits *limited rationality*. There is no point in attempting to build a synthetic cognitive model of someone who truly behaves randomly or in ways that are directly counterproductive to his own intentions. As it happens, history is rich with examples of leaders who were neurotic, or even psychotic in some respects (Hitler and Stalin are examples), but even they exhibited the limited rationality referred to here.

As background, consider that "rational decisionmaking,, as that term is dubiously used by decision theorists, usually assumes: (1) comparing an appropriate set of options by considering the utilities of various outcomes, the probability of those outcomes for each option, and a calculation such as how to maximize expected utility; (2) that the utilities obey the transitivity principle (if A is better than B, and B is better than C, then A must be better than C); (2) a reasonable search for those "appropriate options,,; and (3) the avoidance of various psychological blunders such as seeing only organizationally blessed options, groupthink, sequential rather than parallel assessments, underestimating risks, and so on. None of us, even on the best days, fulfills all these requirements.[10,14,15]

Definition. Limited rationality establishes lower standards. In particular, we defined it as involving:

- A fairly good set of options ("reasonable,, search),
- Parallel examination of options with an eye on objectives, and
- Explicit consideration of *upside and downside potential as well as likely outcomes*.

This does not require formulating anything like utility functions. Decisions made under limited rationality will have a superficially logical basis—i.e., a "reasonable,, relationship between objectives and decisions.[c] They may, however, suffer from a wide range of errors and misperceptions. These include: inappropriate framing, thresholding, with the effect of eliminating what are perceived to be low-probability outcomes; misperceptions; groupthink; etc. Importantly, value judgments are often made implicitly in limited rational decisionmaking. For example, a risk-acceptant decisionmaker may look at likely, worst-case, and best-case outcomes of an option and focus in on the likely and best-case outcomes—unless the worst-case outcome seems obviously too likely to be discounted.[9,10,14,15]

Limited rationality, then, does not imply wisdom or good choices. Rather, it means only that the decisionmaking has satisfied certain minimum criteria for what passes in the real world as rationality. Table 1 illustrates an important generic concept here, that under limited rationality we tend to make decisions by looking at options (one way or another)

---

[c] Models demonstrating limited rationality will satisfy the transitivity principle in any single decision. However, their psychological anchor points may shift over time as a function of events, thereby appearing to violate transitivity. In fact, however, the models are merely being allowed go "change their minds,, about the relative goodness of options as circumstances change their view of the world and their baseline prospects.

and by addressing likely outcome, most favorable outcome, and worst-case outcome, but we do not typically turn everything into utilities and compute some kind of expected utility. Indeed, we think about the values of, e.g., “likely outcome„ in fuzzy terms such as “Good.„

Table 1—A Generic Decision Table Under Limited Rationality

Option	Likely Outcome	Most Favorable Outcome	Worst-Case Outcome	Assessment
Option 1 Option *n*				

3.3 Designing the Model of Red from a Top-Down, Strategic Perspective

Although generalizations here are dangerous, most modeling of Red in other approaches appears to proceed bottom-up. One thinks of influences on Red’s decisions, creates related nodes and arcs, and continues. Sometimes, one introduces hierarchy to reduce clutter, and sometimes one revises to combine concepts in sensible ways, but the standard approach seems to be driven bottom-up rather than by a top-down theory intended to impose structure from the outset. To put the matter differently, design may “emerge„, but it is not built in.

The SCM approach is top-down. A fundamental reason for this is that if the results are to be used in high-level decision support, then Red’s reasoning must be reducible to summary statements drawing upon only a few variables. Real-world decisionmakers, when reviewing their options, are unable to keep large number of variables in mind, much less to discuss them coherently with others. Obviously, however, each and every variable can be questioned, which leads to a “zooming in„, as in “Sir, what do you *mean* that Option 2 looks risky? *Why* do you think of it as risky?„ That question might lead to another explanation with only a few variables, but they would be variables contributing to the perception of risk, rather than to the higher level decision directly. And so on. Again, we see MRMPM at work.

A top-down strategic perspective here means that the top-level variables should be appropriate and natural to decisionmakers. They should be the “right„ variables for the particular decisionmaker. The variables should not be too technical and should not convey the impression of relevance and technical objectivity when they are instead myopic and subject to enormous uncertainties. The importance of this is well illustrated by the actual conversations of President Kennedy and his advisors during the Cuban Missile Crisis.[16]

3.4 Conceiving SCMs to Mitigate Chronic Decision Problems

SCMs should arguably be designed in part to address know decision-support problems. The baseline for most top-level decisionmaking is structured discussion supported by staff papers and electronic viewgraph presentations, but ultimately dependent on the personal power and persuasiveness of the protagonists. This usually works well for the United States, which has a mature National Security Council system that seeks to assure that different perspectives can be heard, debated, and iterated. This mitigates some of the classic cognitive errors celebrated in the theoretical literature. Several problems, however, continue to arise and could be quite serious as the United States develops strategies to counter terrorism, sometimes on the eve of crisis or in the immediate aftermath of crises that might involve mass-casualty and mass-disruption weapons. These problems include:[9,17]

- A practical inability to recognize, honor, and discuss *alternative* concepts of the opponent's objectives and strategy. This <u>tyranny of the best estimate</u> can be due, for example, to imperfect intelligence, the immediacy bias, group think, mirror imaging, and physical exhaustion. Even when alternative views exist, they may not be taken seriously. Those who attempt to express minority views may be accused of being too soft (Adlai Stevenson or George Ball) or too hawkish (those who argue today for action against Iraq). At all levels of organization, there is a high price to pay for pushing a contrary view.

- Failure to acknowledge the inherent difficulty of prediction in complex adaptive systems. It is better to plan and act than to wring one's hands, but naiveté about the quality of predictions can lead to bad plans from which recovery is difficult.
- Related to the above, failure to develop well-hedged contingent strategies. Developing sufficiently adaptive strategies is difficult because of the tyranny of the best estimate and the need to maintain security and discipline within a government when matters are contentious. During the war on Kosovo, even developing ground-force options were prohibited.[18]
- Difficulty in mounting timely decisive action. It is notoriously difficult for democratic governments to take timely and decisive actions that might deter, deflect, or undercut the enemy. A key reason is that the necessity for action is often ambiguous, whereas it is certain that the action will be severely criticized by portions of Congress and the media. Even worse, success may yield nothing dramatically visible (an attack is deterred or diverted), whereas failure could be politically catastrophic.

One purpose of SCM should, then, be to help mitigate these known problems. One way to do so is by constructing *alternative* models of the adversary.

3.5 Alternative Adversary Models

In practice, decisionmaking is often beset by serious uncertainties about the adversary. However, as mentioned above, the processes of government, as well as natural desires to simplify, often lead to giving excessive respect to the best estimate. An obvious solution is to introduce alternative models as a matter of *doctrine* and analysis (thereby taking the onus off presenting the minority view). This can be quite helpful in helping people go "outside the box„ and to see that their baseline view may not be certain. Even two models is far better than one.

But how does one develop such models? The elements of an approach include reading and talking with enough people so that one is exposed to differences of view—not only of anointed experts, but also of semi-outsiders with different perspectives. Another mechanism is human war gaming with participants selected to stimulate debate. A third mechanism is *analyzing* the problem—paying less attention to the presumed dispositions of the actors than to the structural factors at work and the feasible options available.

Such alternative models can be developed. One useful tool for characterizing alternatives is the cognitive map (a variant of influence diagrams). Figure 4 compares cognitive maps for two models of Saddam Hussein as characterized in 1990.[9] Model 1 was the intelligence community's best estimate before August, 1990; Model 2 was an alternative motivated, for example, by taking seriously some of Saddam's more bitter and bombastic speeches, which revealed his grandiose ambitions. Model 2 was also more grandiosely adventurous and risk-taking, which was also consistent with facts. In retrospect, the question is why Model 1 prevailed. The answer is that American officials and analysts projected their own prudence, conservatism and incrementalism onto Saddam.

Alternative models, then, can broaden thinking and generate better-hedged options. However, these advantages disappear if the temptation to "pick one„ is followed. Further, and this is particularly profound for thinking about decision support, real people and governments *change* mindset and reasoning pattern as the result of developments and arguments. Thus, it is a serious error to shift to a single model merely because, at a given point, that model looks better.

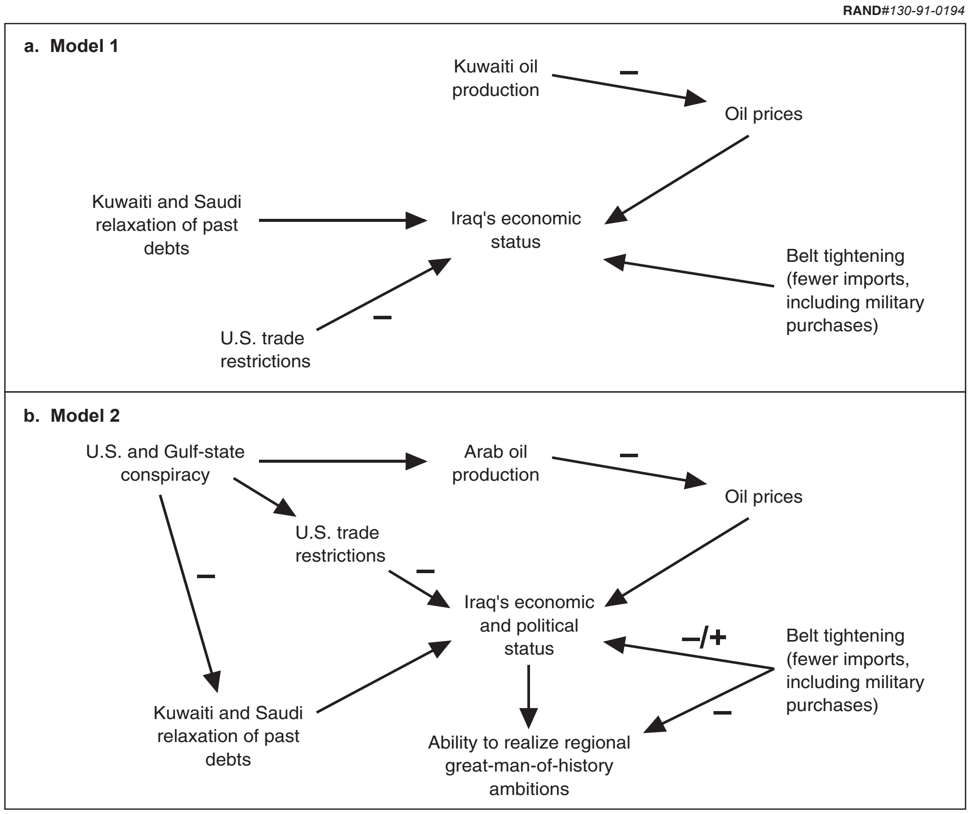

Figure 4—Model 2's Late-July, 1990 Risk-Assessment for the Conquer-Kuwait Option

Table 2 shows another way used to compare the two Saddam models. This one compares their attributes side by side.

The purpose of such devices is to solidify images of the alternative models. Note, however, that this falls short of rigorous definition. To build artificial intelligence models, we would define all of the attributes in terms meaningful to the computer and define what the various attributes would imply about behavior in different circumstances. That is feasible, but tedious. To illustrate the gap between what might look superficially like model definitions (Figure 4 and Table 2) and what we need ultimately, consider how Saddam may have perceived risks when contemplating invasion of Kuwait in mid-1990. How would the Figure 4 and Table 2 map over into such an assessment?

One way to answer that question would be to adopt a Bayesian-net approach and attach numbers, including conditional probabilities, to all of the arrows—with the values parameterized by the attributes. That is a worthy approach and one that has proven its value in a number of applications over the last decade or so. The approach described here, however, takes a very different tack. Figure 5 shows a simple influence diagram from our Saddam work, one suggesting how the second Saddam model would assess risks. This approach hops from a mental understanding of the models sketched above to implications for how the model would think about such matters. In doing so, it does what no before-the-fact computer model can do easily using at-the-time context-specific domain knowledge. Notice reference to "Bush's resolve,,. That might not have been expected to be a key variable if the study had been done in 1988, but—at the time—in mid-1990—President Bush's top advisors were mostly very unenthusiastic about a military confrontation in the Persian Gulf. Even after the invasion, General Powell argued that Kuwait was not worth fighting over.[8] It is easy to understand why Saddam might have doubted Bush's resolve, as well as that of Congress.

Figure 5, of course, is just another picture. How does one move from such a picture to a real model—one that predicts Saddam's judgments and decisions? How does one "add up,, the arrows of the influence diagram, even if it correct? One answer is to use the machinery of Bayesian nets (e.g., by employing MacNetica® or its Windows'® cousin).

Table 2—Comparing Attributes

Attribute	Model 1	Model 2
Ruthless, power-focused; emphasizes *realpolitik*	••	••
Ambitious	••	••
“Responsive„; seeks easy opportunistic gains	••	•
Impatiently goal seeking; likely to seek initiative	•	••
Strategically aggressive with nonincremental attitudes		••
Contemptuous of other Arab leaders	•	••
Contemptuous of U.S. will and staying power		••
Financially strapped and frustrated	••	••
Capable of reversing himself strategically; flexible (not suicidal)	••	••
Clever and calculating (not hip-shooter)	••	•
Pragmatic and once-burned, now cautious	••	
Still risk taking in some situations	•	••
Grandiosely ambitious	•	••
Paranoid tendencies with some basis	•	••
Concerned about reputation and legitimacy in Arab and Islamic worlds	••	
Concerned only about being respected for his power		••
Sensitive to *potential* U.S. power not immediately present	••	•

The approach described here, however, is to construct discrete decision tables such as Table 3. Working from the left, most of the columns correspond to the factors of the influence diagram; the last column is the result. The “values„ of the various factors are qualitative to be in natural language, as in {Very Bad, Bad, Marginal, Good, Very Good}. In constructing such tables, we may wish to study the cases (rows) one-by-one and reason how the model in question would judge risks. This is tedious, but it can be valuable because worrying about the individual cases tends to clarify the factors and their values, and suggest what the attributes of our model should imply. In other words, this discussion forces us to address many of the issues that we might ideally have embedded directly in the model had we gone about things rigorously with the intention of developing an automated artificial-intelligence model. However, the work involved is much simplified doing it specifically for the snapshot situation than in general. Context is concrete and understandable, rather than abstract. Further, we need only to think out those considerations that matter in the context. In practice, this seems not to be especially difficult—although it is a job for a talented interdisciplinary team, not a pure methodologist or mathematician. Lest the reader find this agility dubious, he should note how human war games have consistently proven superior to extensive simulation support over the years, except for “doing the sums„ on logistics, etc. People can be nimble in bringing the right knowledge to bear, even when it requires new thinking. This has also been demonstrated in the Caesar work at war games by Alex Levis and Lee Wagenhals where influence nets have been revised in a matter of minutes to hours.[3]

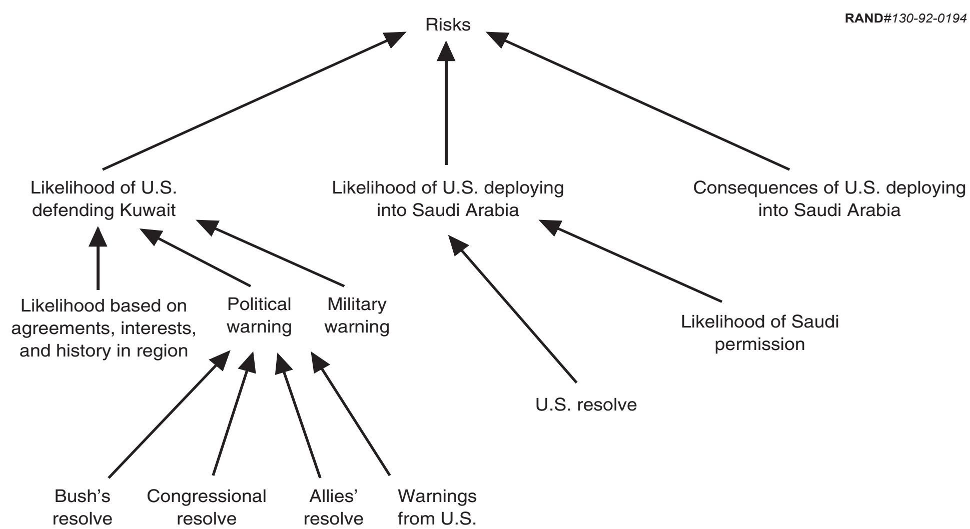

Figure 5—An Influence Diagram for Saddam 2's Assessment of Risks

Although the methodology suggests that one should do some working out of tables such as Table 3 row by row, a more systematic, mathematical procedure would obviously be desirable. In our application to Saddam Hussein, we concluded after deliberation that we could use mathematically simple combining rules *if* we (1) turned the qualitative values into numbers, (2) assumed a linear relationship (Risks= Weighted sum of risk factors), *and* (3) adjusted the weighting factors to reflect the character of the particular model of Saddam under consideration. A risk-taking Saddam would give little weight to worst-case outcomes and would instead focus on the upside. A conqueror or grandiose terrorist leader might be described here as "going for the gold.„ A more conservative Saddam would be very concerned about worst-case prospects, as well as the likely outcome.

Although a nonlinear combining rule might be required in other cases, even these simple linear mathematics can encapsulate a great deal of insight from behavioral psychology and the psychologically sensitive study of past crises. The effect reflects well the celebrated "prospect theory„ introduced by psychologists Daniel Kahnemann and Jacob Tversky.[19] Figure 6 is a schematic representation of that phenomenon. It simply indicates that people are more willing to accept a high ratio of risks to expected gains if their current status (and projected status in the baseline do-nothing option) is wretched than if it is relatively good. This is familiar to all of us. A normal husband and father does not "bet the farm„ on a wager merely because the expected return from the wager would be quite favorable. The decision theorist overly influenced by economist-style decision theory might sneer at this, but common sense recognizes the wisdom. On the other end of things, a sufficiently unhappy person may feel that he has nothing to lose and that only the upside matters. The Palestinians of 2002 are behaving in that manner.

Table 3—A Decision Table for Risks

Likelihood of U.S. Defending Kuwait	Likelihood of U.S. Deploying into Saudi Arabia	Consequences of U.S. Deploying Into Saudi Arabia	Arab attitudes About Invasion	Risks
Low	High	Very Bad	—	Very High
Low	Marginal	Very Bad	Bad	High
Low	Marginal	Very Bad	Marginal or good	High
Low	Low	Very Bad	Bad	Low
Low	Low	Very Bad	Marginal or good	Low
Low	High	Bad	Bad	High
Low	High	Bad	Marginal or good	High
Low	**Marginal**	**Bad**	**Bad or Marginal**	**Marginal**
Low	Marginal	Bad	Good	Marginal
Low	**Low**	**Bad**	**Bad**	**Low**
Low	**Low**	**Bad**	**Marginal or good**	**Low**
Low	High	Marginal	Bad	Marginal
Low	High	Marginal	Marginal or good	Marginal
Low	Marginal	Marginal	Bad	Low
Low	Marginal	Marginal	Marginal or Good	Low
Low	**Low**	**Marginal**	**Bad**	**Very Low**
Low	Low	Marginal	Marginal or Good	Very Low
Low	—	Good or Very Good	—	[Not plausible]
Marginal	—	—	—	High or Very High
High	—	—	—	Very High

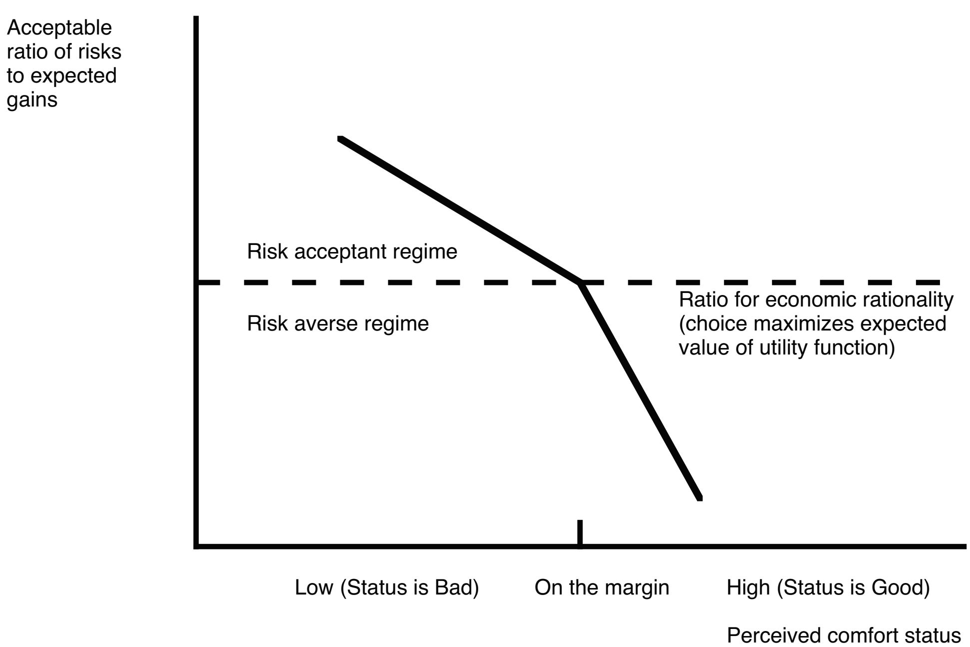

Figure 6—A Depiction of Prospect Theory's Implications

3.6 Extrapolation

This terse account summarizes the basic elements of the methodology for qualitative synthetic cognitive models. One considers snapshots in time at which key decisions will be made, works top down from that decision through the factors that would determine that decision and through subfactors as necessary. For each such decision at a snapshot in time, one constructs the relevant influence diagrams, translates them into hierarchical tables, and populates the tables with what amount to If-then-else rules. The result, however, is not a normal expert system. Instead, it is highly structured, top-down, strategic in nature, and "forward-chaining.„ The reasoning is explicit and easily explainable. It may or may not be automated, but it should be interactive.

In the applications I have conducted with my colleagues, we have found the approach to be relatively straightforward and fast. We did not have to write and debug computer models; we did not have to write down all of the nuances of all of the subtleties. Instead, we were saved by the concreteness of context. We were able to think and analyze quickly. Such terms are relative, however. We could not have done our Saddam work in minutes or hours; instead, we took days and weeks for the high-payoff work and months to tidy things up for the record. Had we had the benefit of knowing the principles beforehand, and of having a variety of computerized tools to help, then it is quite possible that the work could have been done in near-real time—just as it has proved possible in other research projects to adjust belief nets quickly in interactive settings.[3] If we had had the benefit of also working in a decision-support system that allowed virtual conferences with experts of various types at short notice, the quick recovery of historical analogs to the current situation, quick recovery of the adversary's translated and interpreted speeches and writings, ...then a great deal would be possible. Further, the fruits of such work could be translated into structured but comfortable electronic viewgraphs to stimulate and guide high-level discussion. That, at least, is a plausible vision.

Would the results be better or worse than those from some of the other methods being applied? I don't know. I believe, however, that they would hold their own, and then some, with respect to being understandable and readily related to what decisionmakers want to hear about. I believe that their message can readily be translated, for example, into a storytelling form of discussion, or a strong contrasting-view form of discussion, rather than the mysterious calculations of some more technical methods. But a mix of such methods is needed and will prove much better than any one method alone.

A number of improvements in methodology are needed if SCM is to be applied anew. These include: (1) introducing stochastic factors into some of the key judgments and decisions to reflect the fact that, even if one has the adversary model as correct as it is possible to have it, random factors will cause significant variation; (2) adding further rigor and developing the mathematics and computer tools to permit developing and filling out the models quickly; (3) establishing relationships to other methods and suggesting ways in which they can be used to supplement each other (e.g., multiresolution modeling methods could be used routinely in developing belief nets and much of the machinery of systems such as Netica or SIAM could be employed).

4. APPLICATION IN THE WAR AGAINST TERRORISM

An obvious question here is whether these methods can be applied not to Saddam Hussein in a long-ago crisis, but to terrorist groups in the present era. The answer is yes. It is questionable whether it is fruitful to spend much time agonizing about the way Osama Bin Laden thinks. We are unlikely to find ways to dissuade or deter him. By and large, the preferred outcome would be for such terrorist leaders to be killed or captured. This said, the war on Al Qaeda should be seen not as a war on an individual, but rather war on a system, each part of which can be "attacked„ with different instruments, some of them of the sort under study in effects-based planning.[d] Figure 7 illustrates the point by suggesting that the terrorist system consists of top leaders, lieutenants, foot soldiers, logisticians, heads of supportive states, other outsiders (such as wealthy financiers living in third countries), and so on. If the war on terrorism is to be a long affair, then it would be a good investment to develop separate models (with alternatives for each) of each component of the terrorist system. Having such models could, in turn, usefully inform development of effects-based strategies and their adaptation during the course of events.

[d] This discussion draws on work in an ongoing joint project of RAND and IDA, The Institute for Defense Analyses, one sponsored by the Defense Advanced Research Projects Agency. Some of the work has also been reflected in a current study, on counter-terrorism, for the National Academy of Sciences.

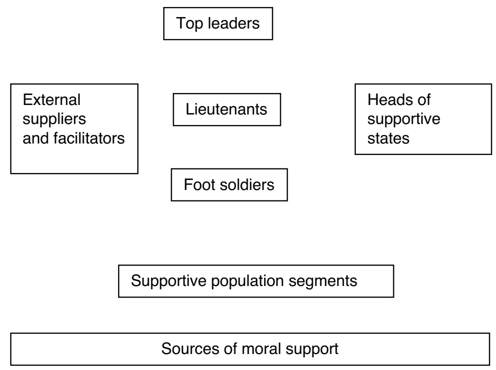

Figure 7—Elements of a Terrorist System

5. ACKNOWLEDGMENTS

This paper is based on RAND research sponsored by the United States Air Force Research Laboratory.

6. REFERENCES

1. Davis, Paul K., *Effects-Based Operations: A Grand Challenge for the Analytical Community*, RAND, Santa Monica, CA, 2001.
2. Paté-Cornell M.E. and S.D. Guikema: "Probabilistic Modeling of Terrorist Threats: A Systems Analysis Approach to Setting Priorities Among Countermeasures,,, submitted for publication in *Military Operations Research*, Working paper of the Department of Management Science and Engineering, Stanford University, 2002.
3. Wagenhals, Lee, "Modeling Effects-based Operations in Support of War Games,,, *Proceedings of the SPIE*, Vol. 4367, ed. by Alex Sisti and Dawn Trevisani, 2001.
4. Davis, Paul K. and James Bigelow, *Experiments in Multiresolution Modeling*, RAND, Santa Monica, CA, 1998.
5. Davis, Paul K., "Applying Artificial Intelligence Techniques to Strategic Level Gaming and Simulation,,, in Maurice S. Elzas (ed.), *Knowledge Based Modelling and Simulation Methodologies*, North Holland Publishing Co., Amsterdam, 1986. (available also as RAND P-7120, 1985).
6. Davis, Paul K., *An Analyst's Primer for the RAND-ABEL Programming Language*, RAND, N-3042-NA, 1990.
7. Davis, Paul K., *Some Lessons Learned from Building Red Agents in the RAND Strategy Assessment System (RSAS)*, RAND, N-3003-OSD. Available also in Military Operations Research Society, *Mini-Symposium Proceedings: Human Behavior and Performance as Essential Ingredients in Realistic Modeling of Combat—MORIMOC II*, 1989.
8. Woodward, Robert, *The Commanders*, Simon and Schuster, NY, NY.
9. Davis, Paul K. and John Arquilla, *Deterring or Coercing Opponents in Crisis: Lessons from the War With Saddam Hussein*, RAND, Santa Monica, CA, 1991
10. Davis, Paul K. and John Arquilla, *Thinking About Opponent Behavior in Crisis and Conflict: a Generic Model for Analysis and Group Discussion,* RAND, N-3322-JS, 1991
11. Davis, Paul K., "Improving Deterrence in the Post-Cold War Era: Some Theory and Implications for Defense Planning,,, Ch. 8 in *New Challenges in Defense Planning*, Paul K. Davis (ed.), RAND, Santa Monica, CA, 1994.
12. Arquilla, John and Paul K. Davis, *Extended Deterrence, Compellence, and the "Old World Order*,,, RAND, N-3248-JS, 1992.
13. Arquilla, John and Paul K. Davis, *Modeling Decisionmaking of Potential Proliferators as Part of Developing Counterproliferation Strategies*, RAND MR-467, 1994.
14. Hogarth, Robin, *Judgement and Choice*, John Wiley and Sons, 1980 (a good summary that includes citations to the seminal research of Kahnemann and Tversky and others).

15. Heuer, Richards J., *Psychology of Intelligence Analysis*, Central Intelligence Agency, Washington, D.C., 1999.

16. May, Ernest R., and Philip D. Zelikow (ed..*), The Kennedy Tapes: Inside the White House During the Cuban Missile Crisis*, Belknap, Harvard University Press, Cambridge, MA, 1997

17. National Research Council, *Post-Cold War Conflict Deterrence*, National Academy Press, Washington, D.C., 1977. (highlights of the author's work are included as Appendix H).

18. Clark, Wesley, *Waging Modern War*, Public Affairs, Perseus Books, Cambridge, MA, 2001.

19. Kahnemann, Daniel and Jacob Tversky, "Prospect Theory: An Analysis of Decision Under Risk,,, *Eonomtrica,* Vol. 47, pp 263-291, 1979.

Thoughts on Higher-Level Adversary Modeling

Paul K. Davis

Reprint from proceedings of SPIE's AeroSense 2003 Conference

Reprinted with permission from *Proceedings of SPIE: Enabling Technologies for Simulation Science VII*, Vol. 5091, April 2003, edited by Alex Sisti and Dawn Trevisani, "Thoughts on Higher-Level Adversary Modeling, Paul K. Davis, pp *172-181*.

Thoughts on Higher-Level Adversary Modeling

Paul K. Davis[a]

RAND and the RAND Graduate School, 1700 Main St., Santa Monica, CA, USA 90407-2138

ABSTRACT

The advent of concepts such as effects-based operations and decision dominance has led to renewed interest in the modeling of adversaries. This think-piece discusses some of the issues involved in conceiving and implementing such models. In particular, it addresses what behaviors may be of interest, how models might be used in high-level decision support, alternative conceptual models, and possible simple implementations. It also touches on issues of multiresolution, multiperspective modeling (MRMPM), modularity, and reusability.

Keywords: Multiresolution, multiperspective, model abstraction, adversary modeling, opponent modeling, decision support, effects based operations.

1. INTRODUCTION

1.1 Objectives

This paper is a think-piece on developing "adversary models" for high-level decision support in political-military contexts. Interest in adversary modeling has increased in modern times as the result of: concepts such as effects-based operations (EBO)[1 2 3] and decision dominance,[4] the DoD having moved to capabilities-based planning,[5 6] and the need to have intelligent behaviors in simulation. A broad theme in planning is achieving adaptiveness, which requires that one anticipate to some degree the *possible* actions of a responsive adversary, rather than only so-called best-estimate projections. A number of related DoD research projects are ongoing, but my purpose here is to start afresh in thinking about some basic issues—while emphasizing *simplicity* appropriate to high-level discussions.

1.2 Approach

In what follows, Section 2 discusses specific behaviors that might be sought from adversary modeling for support of national authorities or commanders (as distinct, say, from incorporating tactical behaviors in entity-level military simulations). As discussed in Section 3, the value of assessments of the adversary in decision support depends on the way in which options are framed for the decisionmaker. I suggest some important departures from nominal military decision-support doctrine. Section 4 then returns to the paper's main flow and discusses some alternative approaches to "thinking about" adversaries and implications for modeling intended to be simple. It also proposes a different way of thinking about model validation. Finally, Section 5 speculates about reasonable and unreasonable goals for composability and reusability of the adversary models discussed. Section 6 summarizes conclusions.

2. WHAT DO WE SEEK FROM ADVERSARY MODELS?

2. 1 Behaviors of Interest

Human behaviors are important at all levels of conflict.[7] At the low end, we may be interested in how effectively pilots function depending on cockpit displays and the volume of information. This paper is concerned more with adversary behaviors relevant to top national authorities, a combatant commander, or major subordinates. Table 1 identifies a number of concrete behaviors relevant to the development of strategy at such levels, strategy such as that for the recent

[a] pdavis@rand.org; phone 310-451-6912.

war with Iraq. The table distinguishes among behaviors of the adversary leader, high-ranking officers, rank-and-file military personnel, ordinary citizens in cities, and zealots—a partial list illustrating the point that "the adversary" is seldom a single entity with a single mindset.[8] If we look first at the political-leader adversary (first column, lines 2 to 4), then for the Iraq war this was Saddam Hussein (if he survived the initial strike). The behaviors of interest were (second column) decisions about escalation, termination, and defensive strategy. The third column suggests some of the possible determinants of behavior. For example, Saddam's decisions about escalation were probably determined by his values (did he care more about putting a thumb in America's eyes than about avoiding widespread death and destruction within Iraq?), how he assessed personal and national prospects for the escalation and no-escalation options, and Iraq's actual military capabilities.

Table 1—Examples of Adversary Behaviors Relevant to Decision Support

Adversary	*Behaviors of Concern*	*Possible Determinants of Behavior*
Political Leader (e.g., Saddam Hussein)	Escalation (none, discriminate tactical WMD, strategic WMD, last-gasp Samson option)	Values, capability, (perceived) personal and national prospects, psychological state
	Termination (capitulation, negotiation, none)	Personal and national prospects, values, psychological state, knowledge
	Defensive strategy (forward defense, defense in depth, fall back to cities)	Prospects for victory, exhaustion of adversary, …, knowledge
Senior Military Officers	Escalation, given orders to do so (as above)	Loyalty, morale, prospects, values, psychological state, quality of command and control system, knowledge
	Defensive strategy (fight forward, defense in depth, fall back to cities)	Prospects for victory, exhaustion of adversary, …, knowledge
	Enthusiasm and speed of maneuver, given orders (low, high)	Loyalty, morale, prospects, values, psychological state, quality of command and control system, knowledge
	Treatment of desertion (acquiesce, discourage, kill)	Loyalty, morale, prospects, values, psychological state. quality of command and control system, knowledge
Rank and File Soldiers	Fighting effectiveness while fighting (low, nominal, high)	Quality of command and control, unit cohesion, loyalty, perceived prospects, physical and psychological state, knowledge
	Break point (low, nominal, high)	Quality of command and control, unit cohesion, loyalty, perceived prospects, physical and psychological state, knowledge
	Termination (surrender, desert, none)	Quality of command and control, unit cohesion, loyalty, perceived prospects, physical and psychological state, knowledge
	Enforcing orders on peers (yes, no)	Quality of command and control, unit cohesion, loyalty, perceived prospects, physical and psychological state, knowledge
People in Urban Areas	Movement (huddle at home, move to safer quarters, flee the city)	Orders, command and control, knowledge, alternatives
	Support of fighters (active, minimal, sabotage)	Loyalty, nationalism, command and control, knowledge, alternatives
	Reaction to Coalition forces (positive, acquiescent, hostile)	Loyalty, nationalism, command and control, knowledge, alternatives
Zealots	Tactics (individual suicide missions,…, wolfpack ambushes…)	Degree of commitment, knowledge, physical and psychological state, support of population
	Termination (desert, none)	Religious zealotry, knowledge, physical and psychological state, support of population

His assessments, of course, reflected perceptions of reality, not reality itself. In the 1991 Gulf war, Saddam apparently did not comprehend beforehand how devastating U.S. airpower would prove to be.[9] Nor did he expect that the United States—the nation embarrassed in both Vietnam and Lebanon—would deploy 500,000 men and mount a ground-force counterattack.[10] Another determinant of decisions is the state of psychological health. Decisionmaking can deteriorate to the point of paralysis when we are ill, exhausted, and frightened.[11]

Skipping in Table 1 to senior military commanders, we have additional determinants such as loyalty. The United States made major efforts to contact Saddam's military commanders and to persuade them to disobey orders and even to surrender.[12] One issue was penetration (e.g., were the leaflets, e-mails, broadcasts, and telephone calls getting through?); another issue was their credibility (and that of the United States).

This partial discussion of Table 1 suffices to demonstrate that the *nature* of behavioral models for supporting high-level decisionmaking is different from that the nature of usual combat or force-planning models, especially in the many qualitative considerations and uncertain values of related variables. To be sure, good combat models also require soft variables and qualitative modeling if they are to be realistic.[7] Unit break points, the impact of surprise on fighting effectiveness, and the force multiplier for attacking unprepared flanks are all examples of such soft variables that can be crucial for, e.g., rapid decisive operations, shock and awe,[14] effects-based operations,[1 2 15] and decision dominance.[4]

2. 2 Value of Adversary Models and the Framing Issues in Decision Support

Given interest in the behaviors discussed in Table 1, a next issue is how related behavior models might actually be used in high-level decision support in crisis or conflict. The context is quite different from, say, an air-component commander worrying about the adversary's tactics in employing surface-to-air-missiles (e.g., employing them all from the outset, versus holding some in covert reserve). It is also different from peacetime force planning that uses simulations evaluating modernization options in scenarios that assume that both sides use their game-theoretic optimal strategies.

As an example of high-level decision support in war, consider that a the commander's choice of maneuver strategy may be affected by his personality, experience, sense of what a"smart" adversary behavior would be, command and control status, and his contextual assessment of the adversary. Robert E. Lee and Stonewall Jackson exploited the inability of early Union commanders to maneuver quickly. Israeli generals in 1967 and 1973 acted upon their belief that Arab armies were unimaginative, slow, and nonadaptive in their maneuvers.

Speculating with such considerations in mind, Table 2 describes schematically some alternative enemy courses of action (COAs) that might be contemplated by a Blue commander on the offensive (such as the U.S. as it reached Baghdad).

Table 2—Illustrative Adversary Defense Strategies (Schematic)

Adversary Strategy	*Static Set-piece*	*Mobile Defense in Depth*	*Area Defense*
Concept	Heroic forward defense	Cause unacceptable high-intensity attrition and hope for negotiation	Preclude quick Blue victory; impose attrition; paint picture of prolonged, bloody war.
Positioning	Forward defense, two up and one back, tactically and operationally	Succession of fixed lines with sizable maneuver forces in reserve.	Relatively small units dispersed throughout area for in-depth attrition and harassment
Maneuver Character	Short-distance, as in reinforcing regiments	For counterattack, flankng, raids, etc.	Short-distance maneuver of small units for ambushing and survival
Use of Terrain	Static defense in open terrain	More mobile operations but still mostly in open terrain	Dispersed defenses in mostly open terrain, plus extensive operations in urban areas
Adaptiveness	Minimal (predetermined behavior)	Substantial	Small units to seek opportunities to attack rear, etc.
Response to Failure	Hold until broken	Continue until defeated	Continue until defeated; shoot deserters; employ irregulars for harassment and suicide missions

Table 3 then indicates how two maneuver options (a deliberate approach versus an approach with more daring deep thrusts, albeit with a fallback option) might be evaluated in the classic doctrinal approach in which the Blue commander first characterizes the goodness of each of his options by assuming that the adversary does the worst that can be done, and then chooses that COA with the least-bad result. This is equivalent to a minimax strategy in classical game theory. As noted in the "story" below Table 3, the commander evaluates his COA for rapid, deep thrusts as being much more risky and therefore rejects it.

Table 3—COA Assessment Following Traditional Minimax Doctrine

Blue COA	*Assessment for Worst Adversary Strategy*
Deliberate attack	Marginal
Immediate daring thrusts (with option to fallback to deliberate attack)	Very Bad

Story: Given a "deliberate" Blue attack, the adversary will surely be defeated, if only by overwhelming power, after Blue tightens lines, builds force strength, and spends time preparing, with air forces and artillery, for further movement. Significant losses will occur with the adversary choosing a mobile or area defense, but losses and other risks should be containable. In contrast, with the daring-thrust approach, it is possible that the adversary—with either a mobile defense or an area defense—could "pull in" and then attack the penetrating forces, especially in urban terrain. And, the adversary could mount major raids into Blue's rear areas. Losses could be high and risks are considerable because Blue's forces are currently strung out and rear areas are vulnerable. Conclusion: The "deliberate attack" is more prudent.

This, however, is not the only way in which options might be assessed and presented. Table 4 provides a very different presentation. In this case, the commander is "encouraged" by the display to think not only of minimaxing, but also about "opportunities." In this depiction, there is a column showing that the best plausible outcome is rated as "Very Good" for the daring-thrust option, substantially better than the "Marginal" outcome of likely and worst cases. The table shows the net assessment to be better for the deliberate attack because it weighs equally the most-likely, best-, and worst-cost outcome cases. That weighting, however, would be context and commander dependent. If presented with this display (or if he demanded its equivalent, regardless of some doctrinal preference for the minimax analysis), an aggressive commander might—whether or not wisely—choose the fast, deep-thrust option. In effect, he would be giving greater weight to the upside opportunity or discounting the risk

Table 4—Assessment Given Ground-Force Focus and Conservative Assessment of Adversary Behavior

Blue COA	*Most Likely Outcome*	*Best Outcome*	*Worst Outcome*	*Net*
Deliberate attack	Marginal	Marginal	Marginal	Marginal
Immediate daring thrusts (with option to fallback to deliberate attack)	Marginal	Very Good	Very Bad	Bad

Story: As discussed above, risks are much higher with the COA of immediate daring thrusts. However, that option also has a substantial upside potential because it may well be that the adversary is fragile, demoralized, and ready to collapse—if only he is engaged quickly and hard. If one weighs the downside risks equally with the upside potential, then—on balance—the deliberate-attack option is still preferable, despite the possibility of losing an opportunity.

Table 5 now shows evaluations if the commander is convinced that the adversary would probably be unable to pull off the more stressing strategies (mobile defense or highly distributed small-unit defenses) effectively, as described by the story underneath the table. In this case, he might reasonably go for the COA of immediate daring thrusts rather than minimaxing.

Table 5—Assessment Given Ground Focus but Discounting of Previous Worst-Case Adversary Behavior

Blue COA	*Most Likely Outcome*	*Best Outcome*	*Worst Outcome*	*Net*
Deliberate attack	Marginal	Marginal	Marginal	Marginal
Immediate daring thrusts (with option to fallback to deliberate attack)	Marginal	Very Good	Bad	Marginal

Story: Because the adversary has never trained in distributed small-unit area defense and has a command and control system unsuited to it, and because the adversary's demonstrated maneuver skills have been only modest in recent days, the worst-case outcomes are unlikely to be terribly bad: the adversary would not be able to prosecute either strategy very well, even if he had the imagination to try. Thus, balancing risks and opportunities equally, the COAs now look about equally attractive, However, the upside is so attractive for the immediate-thrust option that it has the edge unless risk aversion is particularly high.

Finally, consider how the Blue commander's assessment might change as a function of his assessment of how airpower would enter the problem. This is suggested in Table 6. If the Commander had doubts previously, he could now be confident that the more worrisome adversary strategies are infeasible: airpower would devastate forces attempting to maneuver and could protect penetrating units. If the commander also believed that it was *likely*, not just possible, that the adversary's forces would collapse if hit quickly and hard, then he would be even more confident.

Table 6—Assessment Given Discounting of Adversary's Talents and Optimism about Blue's Airpower

Blue COA	*Most Likely Outcome*	*Best Outcome*	*Worst Outcome*	*Net*
Deliberate attack	Marginal	Marginal	Marginal	Marginal
Immediate daring thrusts (with option to fallback to deliberate attack)	Marginal	Very Good	Marginal	Good

Reasoning: As in Table 5, except that the worst outcome for the daring-thrust approach is less bad because airpower will likely be able to prevent enemy maneuver in strength and, if necessary, provide the close support necessary to extricate forward-thrusting units that get into trouble. Thus, because of the very high upside and insignificant downside risk, the immediate daring-thrust option is clearly superior. The worst that is likely to happen is ta fallback to the deliberate strategy, but without having lost much.

The purpose of this notional discussion was to explain how important the framework for COA assessment is and to illustrate how the commander's assessment of the adversary's behavior could enter the problem even in a structured framework. Having illustrated why even simple adversary modeling might be valuable, let us now discuss the nature of useful adversary models for this purpose.

3. ALTERNATIVE IMAGES OF ADVERSARY MODELS

3.1 Conceptual Decision Models

To pursue this papers think-piece style of discussion based on examples, suppose we continue to focus on the top decisionmaker (e.g., a Saddam Hussein) and the decision about whether to escalate. How do we "think about" his reasoning? Figures 1 and 2 indicate schematically a depiction that is perhaps at the right level of detail for high-level discussion. Figure 1 asserts that the likely determinants of his decision are probably his values (e.g., does he value going out with a cataclysm, does he care about the ensuing devastation of his people and country?); his physical capability to escalate; his perceptions about the relative prospects for him personally and perhaps his nation, depending on whether he escalates; and his physical and psychological state (e.g., apocalyptic or paralyzed).

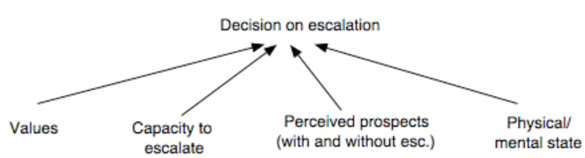

Figure 1—Influences on Escalation

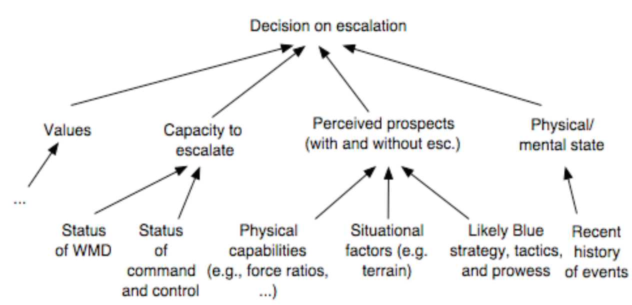

Figure 2—Further Decomposition (Partial)

Figure 2 illustrates how some of the abstractions of Figure 1could be decomposed into more detailed considerations, such as the state of his WMD-related forces and his command and control. This type of influence diagram lends itself readily to hierarchical decomposition using principles of multiresolution, multiperspective modeling.[10 16] Such figures are not mere cartoons for briefing, but rather descriptions of a conceptual model—descriptions largely independent of implementation or representation. Such influence diagrams convey much of what modeling has to offer. Similar breakdowns can be seen, although typically without much thought to the MRMPM theory, in the causal diagrams used by those doing Bayesian-net or influence-net modeling of adversaries.[17 18 19]

3.2 Turning a Concept into a Computational Model

Figures 1 and 2 are conceptual, but we also need a computational model. Many possibilities exist, but the most simple-minded way to proceed (to be criticized later) is with a model such as

$$Prob\_esc = C_1 X_{values} + C_2 X_{cap's} + C_3 X_{benefit} + C_4 X_{mental}$$

where the probability of escalation is approximated as a linear sum of the factors identified in Figure 1, with each X variable being assigned, e.g., a value of 1 to 5, with 5 being most conducive to escalation. This would be a computable model. All that would be required is getting subjective inputs for the X values and the coefficients. That could be done "at the top," without bothering with more detailed analysis, or it could be informed by decompositions, as suggested by Figure 2.

How would such a model be validated? In fact, it seems necessary to reconceive the very concept of model validation because both structure and the values of variables are inherently so uncertain. Retrospective validation might be possible using the historical record and methods used by political scientists,[20 21 22] but that would have limited value for decision support. Community-wide debate would be useful, but it seems that the appropriate way to think about validation in such instances is to ask (1) is the mathematical and logical form of the model reasonable? and (2) is the range of parameter values considered consistent with the state of knowledge?[23]

What, then, can we say about the validity of the model postulated above? A linear model necessarily implies that one can compensate for a shortage in one term by an increase in the size of another term. Is such substitutability appropriate here? The answer is no. To see why, consider a completely different computational representation of the conceptual model suggested by Figures 1 and 2. Table 7 is a decision table of the form used in a good deal of past work—some of it in artificial-intelligence modeling and some of it in more qualitative settings. [10 24] One reads the table as follows (using the second line of the table as an example): If the values favoring escalation are yes *and* the assessment…is yes *and* the mental-health propensity for escalation is yes *Then* the chances of escalation are 0.75, Else..[go to next line of table]. The symbol, --, means that the value in question does not matter. Such tables can be used directly in a high level programming language such as the 1980s'-era RAND-ABEL.[25 26] The table is equivalent to about 50 lines of normal If-Then-Else code, and is far easier to comprehend. Our interest here, however, is in its role as a conceptual model.

Table 7-An Illustrative Decision Table for Escalation

Values favoring Escalation?	*Assessment of improving prospects by escalating*	*Capability to escalate*	*Mental-health propensity to escalate*	*Chances of escalation*
Yes	Yes	Yes	Yes	0.75
Yes	No	Yes	No	0.50
Yes	Yes	Yes	No	0.50
No	Yes	Yes	Yes	0.50
Yes	No	Yes	No	0.25
No	Yes	Yes	No	0.20
No	No	Yes	Yes	0.20
No	No	Yes	No	0.10
- -	- -	No	--	0

This table imagines for simplicity that the variables have values of either Yes or No. Even so, one can quickly peruse the table and see what is going on. What stands out is that if the adversary lacks the capability to escalate, escalation cannot occur (last line). Thus, if this table is roughly correct, the capabilities factor should enter as a multiplier, rather than an independent term, in any equation alternative. The simple linear equation above, then, would be invalid. A second observation is that the table assumes that the "other" factors all have essentially equal effects, except that the "values" factor is a bit more important. And, finally, even if all of the "other factors" are No, the model asserts that there is a non-negligible possibility of escalation (10%).

This, then, is a very different model than the one presented earlier. It still cannot be "validated" in the normal sense, but it is as transparent as I can now imagine (although some people might prefer a tree depiction). Further, creating the table and filling it out is not difficult. The rows are implied by logic and could be generated automatically from Figure 1. The chances of escalation in the last column could be elicited simply by asking, e.g.:

- If all the factors we've identified tilted toward escalation, what probability would you assign to escalation? (Answer shown: 0.75, because the imagined expert doesn't want to imply certainty)
- Are any of the factors *necessary* for escalation? (Answer shown: yes, "capability")
- If the factors identified all argue against escalation occurring, what probability would you assign (Answer shown: 10% because the expert believes that "random factors" can still play a role)

Assuming that the expert sees the table filled in to reflect the previous questions, the last question is

- Can you fill in the "other factors" roughly, using the bounds already set? (Answers shown are values of 0.20, 0.25, and 0.50, chosen to indicate a slight extra significance of "values"," with a yes in column 1 worth slightly more than elsewhere).

The value of this type of approach is considerable. It is easyand—unlike with some of the methods commonly used to elicit information—the underlying model presents the expert with well-defined cases, rather than asking him to answer questions that do not adequately define context.

One further embellishment is suggested by Figure 3. Here inputs or outputs of Table 7 could be specified as probability distributions. For example, the expert being questioned could specify nothing more than the lower bound, most-likely value, and upper-bound. Then, software could automatically generate the appropriate triangular distribution function. With some further assumptions, software could generate a smooth distribution with "tails" added to reflect the well known problem that people tend to have undue confidence about upper and lower bounds.[27]

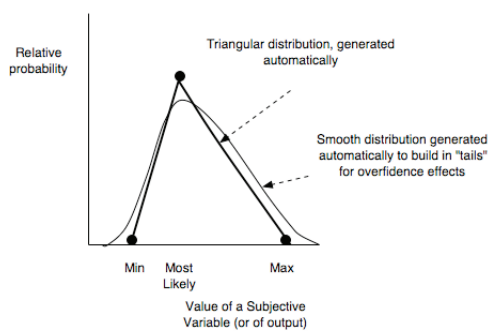

Figure 3—Adding Subjective Probability Information

And, of course, the analyst could readily parameterize the model for sensitivity analysis (e.g., varying the values of 0.75 and 0.1 indicated in Table 7, which were obviously nothing more than rough estimates). If working only with Yes/No values didn't provide enough resolution, then software could generate strawman tables with greater resolution or an equation such as below, with the coefficients determined by normalization and agreement with Table 7.

$$Prob\_esc = X_{cap}[C_1 X_{values} + C_2 X_{benefit} + C_3 X_{mental}]$$

3.3 Summary Speculations on Validation

Although offering no proof, it seems to me likely that:

- Top-level completeness of such a qualitative model (i.e., inclusion of all top-level variables) can be "reasonably" assessed by brainstorming with structures such as Figure 1. In my experience, so long as thought is taken to consult an appropriate range of people, there can be prompt convergence on a set of top-level variables, except that opinions will vary over what names to give them and other details. Still, one might be relatively confident of having identified the key variables (deciding which final set to use, and their names, would be an exercise for the analyst).
- The same will probably not be true at lower levels of hierarchy because participating individuals will have different views of how to decompose the problem. I am skeptical about there existing a "best" ontology—the reason that one will need multiresolution, multiperspective modeling (MRMPM),[23] not just multiple levels of resolution.
- Knowing how the variables of a qualitative model interact is difficult. However, building a simple-table structure, such as Table 7, can quickly illuminate the existence of critical components (such as "capability" in the example above) and suggest related nonlinear forms amounting to interactions.[28] In contrast, simplifying assumptions that assume away interactions are a worry because human behaviors so often depend on combinations of factors: e.g., the combination of, capability, and opportunity. Making such simplifications is not uncommon for people who begin with Bayesian-net approaches because the number of required inputs is so daunting.
- It also seems plausible that experts will be able to suggest some approximate time dependences. For example, immediately after an event such as a shock-and-awe attack, planners might believe an enemy leader to be mentally paralyzed. However, they might reasonably argue—buttressed with empirical data on such matters—that the "relaxation time" for that paralysis would be 12 hours, not 3 days, and that it might be followed by a shift to bitter rage and/or apocalyptic thinking, thus suggesting a window of opportunity. Phenomena such as this have been reported in influence-net work within human war games[18] and are consistent with the way some modern military planners are thinking about effects-based operations within a campaign.
- Finally, it should be possible to map simple and transparent table structures, such as Table 7, into approximate inputs of Bayesian-net analysis. After all, a given line of the table corresponds to specifying a complex conditional probability. It may well be that the theory for accomplishing such things has been worked out already, but in any case populating a Bayesian-net analysis with information gained in such simple and transparent ways might have considerable benefit, especially when working at high levels.

There are no prescriptions for getting structure "right," but methods such as these could lead to models that would be more useful than not, if there were sufficient built-in protections against cognitive biases[29] and sufficiently sophisticated analysts and users at the helm. As emphasized previously, however, it would be essential to use the emerging methods of exploratory analysis to assess the significance of uncertainties.[30] The search would be for relatively robust strategies, not strategies that would be optimal only if the subjective judgments and fuzzy intelligence data were "correct."

4. INITIAL THOUGHTS ABOUT COMPOSABILITY AND REUSABILITY OF ADVERSARY MODELS FOR HIGH-LEVEL DECISION SUPPORT

Given that building useful adversary models should be feasible, it is only natural to ask whether such models should be built with composability in mind. There are excellent reasons for seeking model composability. These include the virtues of modularity in complex designs, the necessity of such modularity when building sufficiently complex systems or systems of systems, the potential reusability of code, which can save large sums of money in some cases, standardization on validated components, and even the stimulus of competition in building the best components.[31]

It may sometimes be logical to seek such composability for adversary models, especially when dealing with large and expensive low-level systems such as those using computer-generated behaviors of military forces described at the entity level in digitized terrain. However, composability and reusability are not likely to be especially useful for the kinds of models discussed in this paper. Instead, the premium should be on simplicity, clarity, and explainability. The reason is that good leaders demand such things in making decisions of great consequence. It is most unlikely that they would find outputs from complex computer models useful unless the results could be convincingly explained in simple terms (i.e., with a simple model or a related and coherent story). So also, the models should use the most meaningful and unpretentious language possible. Still, such models should be "validated," but the best way to do so is likely to be exposure and debate, which again benefits from simplicity of form—e.g., the influence diagrams and tables illustrated above. In candor, I should note that my conclusions stem from the belief that for the highest-level issues, more detailed analytical modeling is unlikely to pay its way because of massive uncertainties that are not reduced with detail.

It also seems likely that the most effective use of high-level adversary models will be in standalone mode using simple data editors and scripts to drive the models with different situational factors. That is, the exploratory analysis can be done with standalone models informed, of course, by detailed intelligence and detailed analysis where relevant. That "informing" could be done by "sneakerware." The high-level adversary models could also be connected to more detailed models, which might be very useful for in-depth background analysis. Doing so is clearly feasible, as my colleagues and I demonstrated in the 1980s with the RAND Strategy Assessment System (RSAS). It is imperative, however, that none of the crucial qualitative variables be dropped merely because there is no easy way to generate them from more detailed models. It will be far better to maintain them as standalone inputs as necessary. To appreciate the danger here, consider that the vast majority of U.S. analyses done before the counteroffensive in 1991 anticipated a difficult battle with thousands of friendly casualties. Only those who insisted on taking "soft factors" seriously—even soft factors such as "fighting quality of the Iraqi army" that could not be calculated from details—got things right.

6. CONCLUSIONS

This speculative think-piece argues that adversary models will be most useful for high-level decision support if decisions are framed in natural ways that highlight both upside potentials and downside risks, rather than in ways that build in a minimax approach. In developing such models, the principal benefits are likely to come from careful development of multiresolution, multiperspective influence diagrams and simple but hierarchical decision tables. The highest-level versions of these can provide a rigorous structure while being comprehensible in strategy-level meetings. Valuable exploratory analysis can be done at deeper levels, with key insights presented in meetings—perhaps with a few examples of drill-down used visually to clarify the logic and build credibility.

Less directly useful will be methods that *embed* the high-level behavior models in highly detailed models of combat, logistics, and intelligence. Such composite models are quite feasible, but they are neither transparent nor easy to reason with. They might, however, prove valuable in background studies of the potential consequences of strategies tentatively based on the simpler reasoning. If the adversary models are simply attachable modules, they can be used both in standalone mode and in connection with the more comprehensive simulations. Key qualitative factors, however, should not be sacrificed in order to have the adversary models fully driven by more detailed models and information systems.

In thinking about issues such as model composability and reusability, it would seem that for higher level decision support it would be more appropriate to worry more about easy communication of the basic concepts, via influence diagrams and tables, than of computer programs. This said, useful tools could and should be developed to aid the simple-model-building activities suggested here *and* to map the results into useful approximate starting data for models of other types, including those based on influence nets and Bayesian nets.

ACKNOWLEGEMENTS

This paper is based on ongoing RAND research sponsored by the United States Air Force Research Laboratory (USAFRL) and the Defense Advanced Research Projects Agency (DARPA). I would also like to thank Jonathan Kulick and Craig Martell, who provided very useful comments on a draft version of the paper.

REFERENCES

1. D. Deptula (Brig. Gen. USAF), "Effects-Based Operations: Change in the Nature of Warfare," http://www.afa.org/media/reports/, 2001.
2. P. K. Davis, *Effects-Based Operations: a Grand Challenge for the Analytical Community*, Santa Monica, CA: RAND, 2001.
3. E. R. Smith, *Effects Based Operations: Applying Network Centric Warfare in Peace, Crisis, and War*, Washington, D.C.: CCRP, Department of Defense, 2003.
4. M. Krause, "Decision Dominance: Exploiting Transformational Asymmetries," *Defense Horizons*, 23, February, 2003.
5. D. Rumsfeld, *Report of the Quadrennial Defense Review*, Washington, D.C.: Department of Defense, 2001.

6. P. K. Davis, *Analytic Architecture for Capabilities-Based Planning, Mission-System Analysis, and Transformation*, Santa Monica, CA: RAND, 2002.
7. Military Operations Society (MORS), *Human behavior and performance as essential ingredients of realistic modeling of combat, morimoc ii, vols. 1-2,* Alexandria, VA: Military Operations Research Society, 1989.
8. P. K. Davis, and B. M. Jenkins, *The Deterring and Influencing Component of Counter Terrorism*, Santa Monica, CA: RAND, 2002.
9. R. A. Pape, *Bombing to Win: Air Power and Coercion in War*, Ithaca, NY: Cornell University Press, 1996.
10. P. K. Davis, and J. Arquilla, *Deterring Or Coercing Opponents in Crisis: Lessons From the War With Saddam Hussein*, Santa Monica, CA: RAND, 1991.
11. A. L. George, "Presidential Decisionmaking in Foreign Policy: the Effective Use of Information and Advice," Boulder, CO: Westview, r1980.
12. R. Atkinson, P. Baker, and T. E. Ricks, "Confused start, decisive end: invasion shaped by miscues, bold risks and unexpected successes," Washington Post, *April 13, page A1*, 2003.
13. U.S. Joint Forces Command, *Millenium Challenge 02*, Suffolk, VA: U.S. Joint Forces Command, 2002.
14. H. K. Ullman, and J. P. Wade, *Shock and Awe: Achieving Rapid Dominance*, Washington, D.C.: National Defense University Press, 1996.
15. D. Saunders-Newton, and A. B. Frank, *Effects-based operations: building the analytic tools,* Defense Horizons, Washington, D.C.: National Defense University, 2002.
16. P. K. Davis, and J. H. Bigelow, *Experiments in Multiresolution Modeling*, Santa Monica, CA: RAND, 1998.
17. J. A. Rosen, and W. L. Smith, "Influence net modeling with causal strengths: an evolutionary approach," Proceedings of the command and control research symposium, Monterey, CA, 1996.
18. L. Wagenhals, I. Shin, and A. Levis, "Executable Models of Influence Nets Using Design/Cpn," 2001.
19. B. Bell, and E. J. Santos, "Making Adversary Decision Modeling Tractable With Intent Inference and Information Fusion," *Proceedings of the 11th Conference on Computer Generated Forces and Behavioral Representation*, Orlando, FL, 2002, pp.
20. R. Axelrod, Ed., *Structure of Decision: the Cognitive Maps of Political Elites,* Princeton, NJ: Princeton University Press, 1976.
21. P. K. Huth, *Extended Deterrence and the Prevention of War*, New Haven, CT: Yale University Press, 1988.
22. P. K. Huth, and B. Russett, "What Makes Deterrence Work? Cases From 1900 to 1980," *World Politics*, 36, no. 4, 1984, pp. 509-510.
23. J. H. Bigelow, and P. K. Davis, *Implications for model validation of multiresolution modeling (MRMPM) and exploratory analysis,* MR-1570-AF, Santa Monica, CA: RAND, 2003.
24. P. K. Davis, "Applying artificial intelligence techniques to strategic level gaming and simulation," in M. S. Elzas, pp.
25. N. Z. Shapiro, H. E. Hall,, and R. H. Anderson, *The RAND-ABEL programming language: reference manual,* N-2367-1-NA, Santa Monica, CA: RAND, 1988.
26. P. K. Davis, *An analyst's primer for the rand-abel programming language,* N-3042-NA, Santa Monica, CA: RAND, 1990.
27. M. H. Bazerman, *Judgment in Managerial Decision Making, 2nd Edition*, New York, NY: Wiley, 1986.
28. P. K. Davis, and J. H. Bigelow, *Motivated Metamodels: Synthesis of Cause-Effect Reasoning and Statistical Metamodeling*, Santa Monica, CA: RAND, 2003.
29. J. Kulick, and P. K. Davis, "Judgmental Biases in Decision Support for Air Operations," *Proceedings of the SPIE*, 5091, Enabling Technologies for Modeling and Simulation VI, edited by Alex Sisti and Dawn Trevisani, 2003, pp. unknowwn.
30. P. K. Davis, J. H. Bigelow, and J. McEver, *Exploratory Analysis and a Case History of Multiresolution, Multiperspective Modeling*, Santa Monica, CA: RAND, 2001.
31. R. H. Anderson, and P. K. Davis, "Composability for DoD Simulations: Why Is it So Hard: What Can be Done?" *presented at DoD Interoperability Conference, Orlando, FL, March 31*, 2003.

Judgmental Biases in Decision Support for Strike Operations

Jonathan D. Kulick and Paul K. Davis

Reprint from proceedings of SPIE's AeroSense 2003 Conference

Reprinted with permission from *Proceedings of SPIE: Enabling Technologies for Simulation Science VII*, Vol. 5091, April 2003, edited by Alex Sisti and Dawn Trevisani, "Judgmental Biases in Decision Support for Strike Operations", Jonathan D. Kulick, Paul K. Davis, pp *260-271*.

Judgmental Biases in Decision Support for Strike Operations

Jonathan D. Kulick,[*a] Paul K. Davis[b]

[a]RAND Graduate School, 1700 Main St., Santa Monica, CA, USA 90407
[b]RAND, 1700 Main St., Santa Monica, CA, USA 90407

ABSTRACT

Human decisionmaking does not typically fit the classical analytic model, and the heuristics employed may yield a variety of biased judgments. These biases are often considered inherently adverse, but may be functional in some cases. Decision support systems can mitigate some biases, but often introduce others. "Debiasing" decision support systems entails designing DSS to address expected biases, and to preclude inducing new ones. High-level C2 decisionmaking processes are poorly understood, but these general principles and lessons learned in other fields are expected to obtain. A notional air campaign illustrates potential biases in a commander's judgment during planning and execution, and the role of debiasing operational DSS.

Keywords: decisionmaking, decision support, cognitive errors, heuristics and biases, command and control, air operations.

1. INTRODUCTION

Decisionmakers operating at high levels of military command face an environment of uncertainty and abstraction different from those at tactical levels, and the challenges facing designers of decision support systems (DSS) for those operations are concomitantly different. In particular, high-level operational commanders must contend with the behavior of adversary leaders, and with inherently contingent secondary and higher-order effects of actions. These commanders (and their staffs) have at their disposal increasingly great volumes of data and tools for managing, presenting, and analyzing them, as well as models and simulations for evaluating courses of action. In this context a DSS is a double-edged sword; designed and employed judiciously it can reinforce the commander's own skilled judgment, relieve him of some cognitive and bookkeeping burdens, and counterbalance his judgmental weaknesses. It may also reinforce or induce some of those weaknesses, if not properly implemented. This paper briefly reviews some issues in the interaction between commanders' decisionmaking processes and decision support systems, and proffers recommendations for the enlightened design of high-level decision support. While the discussion is oriented towards operational command-control (C2) decisionmaking; many elements should be broadly applicable to decisionmaking in other spheres with similar characteristics of abstraction and uncertainty.

It is by now a commonplace notion that human decisionmakers do not typically follow rational actor norms—assessing probabilities and maximizing subjective utility—of classical decision analysis or multiattribute utility theory,[1] except in very simple or contrived circumstances, or for very complex but well characterized strategic planning problems. Whether individually or collectively, they do not naturally approach decision problems in the classical manner and, even when compelled to do so through a structured decision process, do not evince consistent subjective utilities. Furthermore, many important decision problems are inherently too weakly structured to admit to such methods. Studies in cognitive and behavioral psychology describe departures from the normative rational model in how individuals assess and value uncertain outcomes and how experts make decisions in their real environments. The heuristics and biases school accepts the normative value of the rational actor model, and so regards departures from it as errors in decisionmaking. A few researchers have long questioned this principle, noting that there are advantages and sometimes a deep wisdom in decisions that appear narrowly irrational, findings that underpin the burgeoning field of behavioral economics. The naturalistic decisionmaking school takes a bottom-up approach, and considers whether decisionmaking methods are well suited to the tasks at hand and yield profitable outcomes for the

[*] kulick@rand.org

decisionmakers. These theories and their conflicts are well known, but it remains challenging to make effective use of them in designing tools for decision support.

2. DECISIONMAKING

2.1 The rational model, heuristics, and judgmental biases

Contrary to the assumptions of classical microeconomics and decision analysis, people are often poor at estimating probabilities of uncertain events,[2] and inconsistent (with respect to norms such as transitivity of preferences) in making judgments even when given the correct probabilities.[3] Rather than conduct a complex sequence of estimations according to Bayesian probability theory, and a rational process of weighing costs and benefits of the options available according to multiattribute utility theory, people typically cope with uncertainty by applying a small number of heuristic principles (or cognitive shortcuts). While much of the experimental work in this area involves inexperienced subjects in novel settings, the fundamental results have been borne out (often surprisingly and worrisomely, if to a lesser degree) with experts in realistic settings.[4,5,6] Although there is variation within the population and within individuals under different circumstances, these heuristics are commonly enough employed that a whole catalogue of them is well established, as are the systematic and predictable biases in judgment that they produce; of the many taxonomies that have been propounded, we will later employ one that is especially comprehensive and relevant to DSS, and which is derived independently of any particular model of decisionmaking.[7] Note that these are unmotivated biases (people seeing what they *expect* to see), as distinct from motivated biases (seeing what one *wants* to see), socially determined prejudices, or psychopathologies.

There is a considerable literature on military and security decisionmaking, across the spectrum from tactics to grand strategy, and a sizable subset concerned with errors in judgment associated with excessive risk taking, or "military incompetence."[8] The role of heuristics and biases has been explored at the highest levels[9] and, to a lesser extent, at lower levels,[10] but the operational level remains largely unexamined—indeed, little has been written about any aspect of modern operational theory or practice.[11] Judgmental biases are known to be context sensitive, and different biases are observed in different conditions, so we exercise caution in extrapolating from a well studied domain to a largely unknown one. The *overconfidence* bias, as just one example, is widely cited as a cause of poor decisionmaking; this bias is manifest when a decisionmaker's expressed or elicited confidence in the accuracy of his judgments is systematically greater than his success warrants. Empirically, the overconfidence bias is not uniformly observed in many circumstances,[12] but when it is it may deter experts from using decision aids that allow even novices to outperform them.[13] Decisionmaking—by own forces, allies and enemies alike—at multiple levels of command within real-time constraints is also inadequately represented in military simulations; furthermore, achieving this greater realism will require a more profound grasp and implementation of the behavioral decision sciences.[14]

As an example of a heuristic and its attendant biases, judgment of the likelihood or relative frequency of a class of objects or events often reflects how readily examples of that class are conjured up; this is the *availability* heuristic.[15] Readily available instances or images are assumed to represent unbiased estimates of statistical probabilities, even when not germane. For example, the USSR's Cold War assessment of the likelihood of Germany being a renascent military threat to its interests were biased by the availability of World War II images.[16] In assessing an enemy's behavior, a decisionmaker will often rely on the most available model for decisionmaking—his own plans and intentions. Britain based its pre-WWII estimates of the *Luftwaffe's* size on the basis that the "best criteria for judging Germany's rate of expansion were those that governed the rate at which the RAF could itself form efficient units."[17]

As instances of large classes or common events are typically easier to recall than their rarer counterparts, this heuristic often has considerable practical value—indeed, *all* heuristics can enable making reasonable judgments with a minimum of effort, and in many circumstances the resources required to make optimal judgments are not worth the marginal benefits (if any) of greater accuracy; evolutionary psychology argues that heuristics have conferred advantages to those able to make decisions rapidly.[18] (Even so, natural selection yields only *locally* optimized behaviors—i.e., better than your competitors, not the best possible.[19]) Moreover, the speed and efficiency of heuristics should not obscure their sophisticated constituent mental processes, such as pattern matching. Some contend that the "ecological rationality" of heuristics obviates the tradeoff between speed and accuracy—that is, heuristics can be both fast and optimal.[20] A provocative facet of this argument is that "military incompetence" has itself been adaptive, but might no longer be so with historical changes in the nature of warfare.[21]

The term “bias” also carries a pejorative connotation, which is not intended here; we are concerned with *adverse* consequences of judgmental biases—especially as they contribute to poor judgment in unfamiliar decision circumstances. With the availability heuristic, for instance, the ease of recalling examples is not determined solely by their frequency, and the other contributing factors correspond to systematic biases: e.g., *recency* affects retrievability—an air commander whose recent anti-tank sorties have struck at decoys is likely to estimate the incidence of such a deception as higher than another who has never had the experience or, indeed, than he himself will several years hence.

2.2 Decision support systems and judgmental biases

Decision aids can counteract the adverse effects of judgmental biases, by allowing the user to employ heuristics but warning of the likely biases, and by anticipating likely heuristics and providing information that offsets the effects of their use. (We are not concerned here with fully automated decision tools, which are of little utility to high-level C2 decisionmaking.[22]) But DSS designers (and users, no less) must recognize that decision aids can themselves introduce biases into decisions that might otherwise not suffer from them. For example, the manner in which decision problems are *framed*, such as whether outcomes are represented as gains or losses, influences the choices that are made.[23] Prospect theory holds that decisionmakers tend to be risk averse with respect to gains but risk loving with respect to losses;[24] it is a powerful framework for explaining risky high-level operational[25] and national strategic decisions.[26] Decision support systems that frame options for the user or even present, say, a neutrally phrased checklist for his consideration may thereby bias decisions, even if no weights are implied. Decision aids that incorporate user-driven database or knowledgebase searches may reinforce *confirmation* biases, which stem from a decisionmaker’s tendency to search for information that supports a pre-established hypothesis,[27] and to ignore rebutting information that may arise.[28] Some maintain that senior military commanders and politico-military leaders may be especially prone to such judgmental biases, due to selection bias in intellectual characteristics and to organizational forces,[29] but this view runs counter to more contemporary findings.[30] Military course-of-action (COA) analysis that characterizes a COA by the nominal outcome expected if the enemy takes the worst action possible against it tilts the problem toward a “minimax” style of thinking, which is unsatisfying for an aggressive commander interested primarily in how to achieve his own ambitious objectives. A more balanced approach is to characterize an option by its likely, best-case, and worst-case outcomes,[31] and to then identify the circumstances that would enhance likelihood of the best-case outcome and reduce likelihood of the worst case.

Without empirical findings that a particular bias is manifest in a particular decision environment, its significance should not be assumed. More than that, attempting to infer significance from broader experimental research can be fruitless. A cautionary tale: a casual reading of the literature might suggest that the inadvertent introduction, from using DSS, of one widely observed bias—base-rate neglect—may be quite easily remedied, requiring no reconception of the decision aiding principle, but a more considered assessment leaves the matter unclear. Many judgments require the decisionmaker to combine information about a more or less stable average incidence of some class of events (the “base rate”) with specific information about a member of that class; a commander might know, for instance, that an enemy has only rarely been found to collocate military communications operations in hospitals, and he has fresh intelligence that encrypted radio transmissions are issuing from a particular hospital. How should these transmissions be interpreted in light of the base-rate knowledge? Does the evidence compel an alternative explanation?

Classical decision theory dictates, by Bayes’ law, that prior probabilities inform the interpretation of new information, but many studies have shown that even experts given familiar problems are not intuitively Bayesian. For example, a famous medical school experiment found that even experienced health professionals are not Bayesian in their interpretation of simple test results: told that a disease is present in .1 percent of the population and that the probability of a false positive result on a test for the disease is .05, nearly half of the subjects estimated that a randomly selected person who tests positive has a 95 percent chance of having the disease, and fewer than 20 percent offered an estimate in the neighborhood of the correct value of 2 percent.[32] Now, neuroscience suggests that people have a collection of mental modules for various cognitive functions, including one for reasoning, that perform better with information in formats that humans adapted to. Humans have had considerable direct experience with “natural” frequencies (“the river flooded 4 years out of the last 20”), whereas probabilities are considerably more abstract (“the probability of the river flooding in any year is 0.2”).[33] The frequency formulation, in this case at least, contains more contextualizing information, in the numerator and denominator. When the medical school experiment was replicated some years later, the control group presented with the probabilistic problem formulation performed as poorly as in the original experiment, but the subjects given a frequentist formulation did not ignore the Bayesian prior, and estimated the result correctly.[34]

This is not merely a matter of semantics, of presenting information clearly, as there is nothing ambiguous about the probability formulation; the frequency formulation—its proponents argue—is simply better matched to our basic cognitive processes. But these findings are themselves hotly disputed, on methodological and theoretical grounds.[35] What is more, these countervailing studies contend that decisionmakers presented with frequentist data ignore diagnostic evidence and *overweight* the base-rate. A frequentist formulation for an a priori likelihood estimate under uncertainty (e.g., "I'd say there's a 20 percent chance of finding chemical weapons in that bunker" conveyed as "you'd expect to find chemical weapons in a bunker of that sort two times out of ten") can improperly suggest an empirical-frequency basis for the estimate. Some recent research suggests which statistical format is preferred in different situations,[36] and others contend that an odds-ratio formulation (i.e., "the odds are one to four against finding chemical weapons in that bunker") is the most readily understood.[37] Still other investigators maintain that biases that appear to reflect insufficient reaction to new evidence may serve the decisionmaker well in the face of possible real-world changes that affect the reliability of evidence or its significance.[38] In short, it appears that framing likelihoods as probabilities or frequencies *might* influence decisionmaker judgments, and that those judgments are not robustly Bayesian, but more research is required to derive useful prescriptions for decision support.

For so simple a problem, of course, a DSS could perform the Bayesian calculation and derive the correct result, but if a decisionmaker is to act on it, he should have the result presented in the most meaningful format. High-level decision-aiding modeling tools (such as CAESAR II/COA[39]) often present their results in probabilistic formats, whether as point estimates with standard errors, uncertainty intervals, or probability distribution functions. Many consumers of these results may grasp them intellectually, but still draw biased inferences that could be allayed with a frequency format. Indeed, some other biases may disappear when the statement formulation is changed.

With the *subset* bias (sometimes denoted "conjunction fallacy"), people estimate the likelihood of an object having two independent properties as higher than it's having at least one of them. A common illustration is the "Linda problem": several of Linda's attributes are given, and subjects are asked to rank order statements about Linda according to their probability. Most subjects rank "feminist and bank teller" higher than "bank teller," even when it is clarified that the latter statement does not require concluding that she is *not* a feminist.[40] When the question is given a frequency formulation—"100 people fit the description above; how many are bank tellers, how many are feminist bank tellers,…"—the subset bias is much more weakly observed.[41] Explanations for the subset bias vary. Specific scenarios may seem more likely than general ones because they better represent the way we imagine events.[42] More particularly, the conjunctive label may provide a more compelling pattern match to a real person than the disjunctive labels, which yields the overestimation in the original experiment,[43] while subjects thinking about a large sample are less likely to employ a pattern matching heuristic in the frequentist experiment. Alternatively, people might judge the conjunctive label to be slightly less probable but much more informative than the disjunctive labels, and so to have a higher "expected informativeness."[44] Although the military decisionmaking literature does not address the subset bias explicitly, the possible relevance to operational decisionmaking—and an opportunity or pitfall for DSS—is clear: a commander might imagine, for example, a greater likelihood of an enemy regime collapsing due to an attack on its C2 assets than of the regime collapsing (with causes not further specified).

A prominent recent event highlights the potentially severe consequences of judgmental biases, with special relevance to military operations. In testimony before the independent panel on the space shuttle Columbia breakup, an expert advisor to NASA said that it had

> again fallen prey to "systemic" flaws in reasoning—such as the creeping acceptance of poorly understood risks in operating the space shuttle….[D]espite prodigious efforts and the best of intentions, [NASA] had failed to upgrade its aged database and computer systems to allow it to track subtle but unacceptable trends….[T]he shuttle team has been lulled by repeated successes. "I think there's a flaw in the reasoning of many well-intentioned people" in forgetting that "if you've a 1 in 100 chance of risk of an event occurring, the event can occur on the first or the last [opportunity], and there's an equal probability each time."….[T]he perception within the agency seemed to be "that if I've flown 20 times, the risk is less than if I've just flown once. And we

were continually attempting to inform them that unless they've changed the risk positively, they still have the same issue even after 50 flights or 60 flights."[45]

Several judgmental biases might have been at play here, most notably the overconfidence bias and the disjunction bias, which holds that probability is often underestimated in compound disjunctive problems.[46] NASA engineers and military commanders are equally highly trained, disciplined, and responsible, and yet both are subject to a biasing organizational dynamic—the pressure to weight observed successes more heavily than is warranted. Imagine a commander deploying a innovative new platform in combat for the first time. Development testing suggests that it should fail in about five to ten percent of its sorties against enemy surface-to-air missiles. It survives its first three combat sorties; should this success embolden the commander to rely on it more than he had been inclined to before the first flight? Certainly not—but his inclination to do so should not be surprising. It is no challenge for a DSS to make the apposite calculations, but to convey the results in a compelling manner is nontrivial.

2.3 The naturalistic model

The heuristics and biases school has come in for some serious criticisms, on several grounds.[47,48] As mentioned above and discussed at greater length below, heuristics often yield cost-effective decisions compared with expensive (in time and mental energy) rational processes. Moreover, it may not be worth even a modest effort at an optimizing judgment at any particular time, for dynamic problem situations that will soon obsolete the judgment. Some other criticisms concern the research methodology—that researchers demonstrate selection bias; that they focus on the statistical significance of biases of small magnitude; that they use contrived problems in which one interpretation is deemed normatively correct, ignoring alternatives in which responses seen as reasonable; that they elicit one-off judgments of static problem settings; and so on.. More fundamentally, some critics argue that the normative standard of rationality is itself spurious, so that departures from that norm are not cause for concern if the judgmental biases yield outcomes that their bearers are happy with.

The richly diverse naturalistic decisionmaking (NDM) theories, which have gained great prominence in recent years, focus on how people make decisions in their natural environments—that is, they take a descriptive rather than a classically normative approach.[49] They tend to hold that (1) situation assessment is more important than option generation; and (2) options are considered sequentially rather than simultaneously; (3) are evaluated by mental simulation or pattern recognition; and (4) are accepted if satisfactory, rather than optimal. In broad strokes, they tend to find that people are effective decisionmakers—that heuristics *work*. While qualitative naturalistic models less readily lend themselves to computational DSS than do quantitative classical models, the principle that decision aids should not replace the user's natural approach to decisionmaking or force him into ill fitting rational processes has found some traction in tactical DSS.[50,51]

The strong case of NDM theory does not simply dismiss biases as insignificant or tolerate them as unavoidable side effects of otherwise valuable heuristics. Rather, it celebrates biases as adaptive and situation-appropriate, as does the history of scientific progress, writ large. Scientists form hypotheses—often just glorified hunches—whose proof they pursue vigorously. If the evidence is lacking or disconfirming, they typically adapt the hypothesis and tack a revised course, without dwelling on the prior mismatch between theory and data. Scientific inquiry is, in this sense, descriptive. The strong case of classical decision theory, which damns all biases as defects in decisionmaking, is analogous to a fundamentalist statistical approach to scientific inquiry, which argues that the data should speak for themselves, and against hypotheses generation and testing. Data-generated scientific discoveries, however, have not been nearly so significant as hypothesis-generated theories, and tend to lack much explanatory power beyond the scope of the data.

A full treatment of these controversies is beyond the scope of this paper; they are likely overstated by the more doctrinaire factions in each camp, and a synthesis is possible.[52] Classical models of decisionmaking leave little room for broad general knowledge and more contextual tacit knowledge, and so deprive the decisionmaker (and user of a classically based DSS) of many of the benefits of experience and learning. Naturalistic models, zealously adhered to, fall victim to false pattern matching and willful deception. We stress only that the heuristics and biases findings can expose intellectual limitations and suggest how to improve the quality of thinking, and can reveal processes that guide judgment and inference. And, furthermore, the heuristics and biases school is relevant to DSS, even if the supporting experiments are contrived, because interactions with DSS in real decision environments can be *similar* to contrived experiments in some respects.

3. SUPPORT FOR OPERATIONAL DECISIONMAKING

3.1 Decisionmaking at different levels of command

Empirical research on military decisionmaking has focused almost exclusively on tactical actors and situations, up to the division command level.[53] An extensive research program on naval tactical decisionmaking, spawned by the USS Vincennes incident, has contributed to understanding the implications for effective tactical decision support.[54] There has been any number of retrospective studies of military operations with detailed accounts of high-level decisions[55] and memoirs from top decisionmakers, but no similarly rigorous observational studies of high-level operational decisionmaking. General officers' time is dear, and they are not likely to be available for laboratory experiments; during actual combat operations, likewise, decisionmaking researchers are not given full access to the operations center. (Experts are difficult to study, in general: in addition to the problem of access, researchers face the prospect of becoming well enough versed in the expert's field to be able to judge their performance.[56]) War games could be designed to serve decisionmaking research without compromising their principal objectives, but there has been little rigorous observation on judgmental biases in operational-level wargames; these biases have been identified and studied in lower-level wargames.[57] There are also, by the same token, limited efforts at modeling high-level operational decisionmakers for use in simulations;[58] judgmental biases may also enter importantly into enemy decisionmaking and should be included in model representations of enemies.[59] An excellent high-level study of commanders' information needs focuses on the flow of information between commanders and subordinates in C2 decisionmaking, and mentions in passing the role of judgmental biases (especially overconfidence).[60]

3.2 Decision support systems: debiasing and biasing

More generally (i.e., not limited to military applications), while many DSS are intended to mitigate the effects of judgmental biases, there has been little consideration of how their use may *contribute* to biased decisionmaking (a promising study of biases in the judgment of medical patients using DSS is underway.[61]) A consideration of judgmental biases is not evident in many discussions of military operational DSS, even for those that account for a variety of users' idiosyncrasies. For instance, the Attack Operations Decision Aid (AODA)[62] is a tool to assist in diverting air assets, from missions already specified in an air tasking order, to time-critical targets.[63] While the tool would be forward deployed on airborne platforms and used for routine targeting decisions, very high value, high risk, or politically sensitive targets would bring the Joint Force Air Component Commander (JFACC) or higher-level commanders into the decision process. The tool both supplements and supplants the commander's own decisionmaking capabilities:

> AODA's algorithms are based on operations research techniques. A commander makes similar decisions using a heuristic approach which, while adequate in a non-stressing (few-on-few) environment, can not efficiently handle complex situations. AODA assesses the tradeoffs among original target value and new target value, available weapon capability, asset survivability, and probability of destruction. AODA then provides the operator with a list of recommended weapon target pairings.
>
> AODA's algorithms require that the values of targets and assets be captured numerically. Although subjective valuations of this nature are made by commanders during the decision process, they are not quantified to the level required by the decision aid. To fully support the aid's algorithms, commanders will have to explicitly state the values they place on targets and assets.[64]

These sorts of subjective valuations (even by experts) are prone to judgmental biases across a broad range of fields. It is not evident whether these potential biases are explicitly considered in the tool's design, although they should certainly figure in many of the other individual-level factors that influence in its effectiveness:

> The decision aid must meet the user's perceived needs and incorporate those factors that the user feels are critical to a correct decision. It is imperative to determine what the user thinks the decision support needs are, the conditions under which the aid is needed, the features that are needed, and the factors that the

aid's algorithms should consider. A decision aid should be built with a clear understanding of the users' expectations and level of expertise as well as the operating environment. Decision aids may support various levels of command. At low levels, decision aids may simply help the operators to recognize a critical situation, so that pre-planned appropriate action can be taken and important information can be elevated to other command levels.

Decision aids that support a commander responsible for the execution of the campaign plan, may need to gather all the available data, organize and present information clearly, and recommend options that facilitate decision making. A decision aid's support level of sophistication needs to be geared to the user's training, educational level, and background, which varies with the command level.[65]

With regard to the lattermost point, higher levels of command historically have favored analytical decisionmaking over naturalistic, which should also influence the nature and sophistication of the DSS, but ongoing advances in information technology are blurring the distinctions among the levels and are causing them to merge.[66]

The two senses of reducing bias in using DSS—correcting for preexisting bias and not inducing new bias—go under the rubric of "debiasing." Debiasing also includes training and conditioning the decisionmaker to reduce his propensity to judgmental biases even without the use of aids to particular decisions,[67] which purpose wargames may serve.[68] Bias in decisionmaking can stem from the decisionmaker, the decision environment, or a mismatch between them; some critics of the heuristics and biases approach maintain that experiments that find bias often suffer from such a mismatch, and should themselves be debiased—in fact, the most significant work on debiasing strategies has come from experimental psychologists conducting laboratory studies.[69] In any event, many real-life operational decision environments *are* "artificial" in the sense of laboratory experiments—the problems are unfamiliar, ambiguously defined, and present conflicting goals, and it is not evident that "life is more charitable to people than are experimenters."[70] In many situations DSS may be able to debias the decision environment: making it easier to execute a given process, facilitating the use of a better process already in the decisionmaker's repertoire, or providing an information structure that works better with the process already in use.[71] A number of experimental debiasing systems are described in the literature but none appear to be in active wide use.[72]

4. DEBIASING AN AIR CAMPAIGN

Up to this point, the paper has provided a survey of relevant decision science research. Let us now consider a concrete illustration of the role of judgmental biases in operational DSS: a notional campaign and a commander charged with operational decisions, in this case the JFACC producing the master air attack plan (MAAP) and daily air tasking orders (ATO).[73] We pose plausible circumstances for representative judgmental biases, within a narrative of this campaign and tasks, and consider the possible role of DSS; this thought experiment does not reflect any actual DSS in current use or development, many of which are no doubt well ahead of our thinking in these regards, and entails a caricature of the JFACC's proneness to biased judgment.

The taxonomy we use includes 37 different biases, classified into 6 categories.[74] The narrative will illustrate one bias from each category:

- *Memory* biases: most fundamental, concern storage and recall of information
- *Statistical* biases: non-probabilistic information processing
- *Confidence* biases: act to raise confidence in own judgment and decisionmaking skill
- *Adjustment* biases: undue attachment to initial assumptions
- *Presentation* biases: concern the way information is perceived and initially processed
- *Situation* biases: highest level of abstraction; concern response to the general decision situation

Red has invaded and occupied two zones of its neighbor, Green. Another neighbor, Yellow, is covertly providing support and shelter to Red leadership. Blue is mounting an air campaign to compel Red to withdraw from Green, to deny it the capability to attack its other neighbors, and to prevent it from transferring C2 capabilities or materiel to Yellow. The illustrative biases are:

- *Habit* bias (situation): A Bayesian net model for inferring Red command leadership intent requires the air operations center staff to enter almost a hundred subjective probabilities about Red response to stimuli. In a previous campaign (against a much different enemy), the (Blue) JFACC had a successful experience with the same model, in which a value of 0.2 had been entered for all the probabilities, so he instructs the model operator to do the same in this case. Habit is an extreme manifestation of bounded rationality—choosing an action because it has been used previously.
- *Regression* bias (adjustment): Development testing suggests that a newly deployed bomb will hit within 5 m of the aimpoint, on average, 85 percent of the time; it is configured to be carried by two different aircraft types, each carrying one bomb, with equal accuracy expected from each. On day one, aircraft type A delivers 100 bombs, with 80 hitting within 5 m; type B delivers 200 bombs, with 180 direct hits. The next day's targets will require an estimated 90 direct hits. Impressed with the bomb's performance when delivered by type B, the JFACC dispatches 100 sorties of the second aircraft type, expecting a 90 percent strike rate. He has ignored the likely regression to the mean—if the aircraft are equally accurate, on average, then the type that performed better the first day will not do so consistently.
- *Completeness* bias (confidence): A campaign model provides a prediction of Blue aircraft day one losses, for three candidate master air attack plans; it assumes canonical values for Red air defense capabilities, based on the models and age of their weapons, although the model is capable of higher resolution estimates with inputs on manpower and weapons maintenance. The model outputs best estimates of 3.04, 3.41, and 2.93 losses, respectively, with 90 percent confidence intervals of ± .16, .22, and .16. The JFACC perceives these apparently precise estimates as definitive, and curtails the search for more data to inform the decision. An apparently complete set of data inspires undue faith in the quality of the inputs and assumptions that yielded them. Had the display read ~3, ~3½, and ~3, the JFACC would probably have sought additional input data for higher resolution calculations.
- *Framing* bias (presentation): A Monte Carlo evaluation of a campaign model compares two MAAPs, each using 100 ground attack aircraft; for the first plan, the model predicts 95 aircraft surviving day one, 85 surviving day two, and 70 surviving day three; for the second plan, 100 day one, 90 day two, 60 day three. The JFACC chooses the first option. Prospect theory suggests that he is risk averse with respect to gains (survival rates), and risk seeking with respect to losses; if the outcomes were expressed as losses (fatality rates), he would likely choose the second MAAP.
- *Hindsight* bias (memory): On day six Blue begins to attack fixed ground targets in one occupied zone of Green, in an effort to compel the occupying Red forces to leave, either of their own accord or under orders from higher-level Red leaders. The ATO calls for a total of 24 500-lb precision guided bombs to be dropped on 18 different targets. After dropping one bomb on a munitions depot in an abandoned village, the occupying forces retreat in haste from the entire district, leaving behind their artillery. Pleased with the effects achieved with a single well placed bomb, the JFACC is confident that he predicted this outcome, and that it could hardly have turned out otherwise. He revises the next day's ATO for attacking the other occupied zone, without seeking more information on *why* Red forces fled.
- *Base-rate* bias (statistical): On day eight the JFACC receives credible intelligence that three men in tan uniforms are in a white Jeep with a black roof, on the highway heading to the border with Yellow. A knowledge-based DSS gives a high likelihood that three wanted Red officials fit the description in the intelligence; the commander dispatches a missile equipped drone to find and destroy the vehicle and gives firing authority to the drone operator. The JFACC has ignored (or not sought out) the base-rate data—that *most* of the cars in the area match the description in the intelligence.

In these examples, DSS are explicitly implicated in the habit, completeness, framing, and base-rate biases—the format of the DSS output, the user interface, or the mere fact of employing the DSS stimulates or amplifies the JFACC's propensity to judgmental biases, none of which are clearly benign. More careful design of the DSS could mitigate some of these suboptimal judgments, without imposing undue hardships on the JFACC's own decisionmaking style. In the regression and hindsight bias cases, the JFACC draws possibly biased inferences from valid statistical data and recent observations; in the former case a DSS that monitors the data being collected could generate a warning not to misinterpret short-term deviations from average performance. The latter case presents a thornier problem, as it's not a matter simply of appropriate data display formats or monitoring calculations: a debiasing DSS would have to force the JFACC to consider alternative explanations for what he observed; various strategies of this sort have been found to at best reduce hindsight bias, and recent studies suggest that they can backfire and *reinforce* biased judgments.[75]

5. CONCLUSIONS

Decisionmakers of all stripes are subject to judgmental biases, and senior military commanders are likely to be so, as well. Emerging changes in operational doctrines, driven by advances in technology and shifting political environments, dictate new challenges for decision support and a greater attention to the role of judgmental biases. Even if decisionmakers are made cognizant of, and are trained to overcome, their propensity to biased judgments, DSS can relieve them of some of the burdens of judgment. At the same time, reliance on DSS can evoke biases that might not be manifest in unaided decisionmaking.

It is incumbent upon DSS developers to consider both sorts of biases. A DSS should systematically address all of the known biases, with respect to the particular attributes of the intended users, their decision environments (to include staff processes, staff training, etc.), and the spectrum of decision types. This deliberation should be regarded as an integral part of "validating" the DSS.

The enduring and fruitful tension between analytical and naturalistic standards of decisionmaking suggests that DSS in development should be reviewed by interdisciplinary teams with a diversity of agendas. Behavioral scientists, for instance, might worry about biases from a "let the data speak for themselves" perspective, while former commanders might be more concerned with discounting "suboptimal" opportunities and otherwise being overly analytic. An especial challenge to DSS design is to make available to users more qualitative data and what otherwise would be tacit knowledge, such as "although no hard data are available, the sense of our intelligence is that enemy units are demoralized and unstable, and might collapse if stressed more." This tacit knowledge provides invaluable context for interpreting and acting on conventional DSS displays, such as force ratios and assumed breakpoints.

ACKNOWLEDGMENTS

This work was sponsored by the United States Air Force Research Laboratory, as part of a study by RAND's Project Air Force, a federally funded research and development center. The work also draws on research being accomplished in a study of high-level decision support by the Defense Advanced Research Projects Agency (DARPA) as part of its Genoa II program.

REFERENCES

1. R. Keeney and H. Raiffa, *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*, John Wiley & Sons, New York, 1976.
2. A. Tversky and D. Kahneman, "Judgment under uncertainty: Heuristics and biases," *Science*, **185**, 1124–1131, 1974.
3. A. Tversky and D. Kahneman, "The framing of decisions and the psychology of choice," *Science*, **211**, 453–458, 1981.
4. C. Camerer, "Individual decision making," in J. Kagel and A.E. Roth, eds., *Handbook of Experimental Economics*, Princeton University Press, Princeton, 587–703, 1995.
5. R.J. Heuer, Jr., *Psychology of Intelligence Analysis*, Central Intelligence Agency, Langley, VA, 1999.
6. G.P. Hodgkinson, "Breaking the frame: An analysis of strategic cognition and decision making under uncertainty," *Strategic Management Journal*, **20**, 977–985, 1999.
7. D. Arnott, "A taxonomy of decision biases," Technical Report 1/98, School of Information Management & Systems, Monash University, Melbourne, 1998.
8. N. Dixon, *On the Psychology of Military Incompetence*, Basic Books, New York, 1976.
9. R. Jervis, R.N. Lebow, and J.G. Stein, *Psychology and Deterrence*, The Johns Hopkins University Press, Baltimore, 1985.

10. M. St. John, J. Callan, and S. Proctor, "Tactical decision-making under uncertainty: Experiments I and II," Technical Report 1821, SPAWAR Systems Center, San Diego, 2000.
11. J.A. Warden III, *The Air Campaign: Planning for Combat*, Pergamon Brassey's, Washington, 1989.
12. G. Gigerenzer, "The bounded rationality of probabilistic mental models," in K.I. Manktelow and D.E. Over, eds., *Rationality*, Routledge, London, 284–313, 1993.
13. S.E. Kaplan, J.H. Reneau, and S. Whitecotton, "The effects of predictive ability information, locus of control, and decision maker involvement on decision aid reliance," *Journal of Behavioral Decision Making*, **14**, 35–50, 2001.
14. R.W. Pew and A.S. Mavor, eds., *Modeling Human and Organizational Behavior: Application to Military Simulations*, National Academy Press, Washington, 1998.
15. A. Tversky and D. Kahneman, "Availability: A heuristic for judging frequency and probability," *Cognitive Psychology*, **5**, 207–232, 1973.
16. R.J. Heuer, Jr., "Strategic deception and counterdeception: A cognitive process approach," *International Studies Quarterly*, **25**, 294–327, 1981.
17. F.H. Hinsley, E.E. Thomas, C.F. Ransom, and R.C. Knight, *British Intelligence in the Second World War*, vol. 1, Cambridge University Press, New York, 1979.
18. L. Cosmides and J. Tooby, "Are humans good intuitive statisticians after all? Rethinking some conclusions from the literature on judgment under uncertainty," *Cognition*, **58**, 1–73, 1996.
19. H.A. Simon, "Rational choice and the structure of the environment," *Psychological Review*, **63**, 129–138, 1956.
20. G. Gigerenzer, P.M. Todd, and ABC Research Group, *Simple Heuristics That Make Us Smart*, Oxford University Press, New York, 1999.
21. D.D.P Johnson, R.W. Wrangham, and S.P. Rosen, "Is military incompetence adaptive? An empirical test with risk-taking behaviour in modern warfare," *Evolution and Human Behavior*, **23**, 245–264, 2002.
22. C.D. Wickens and J. Hollands, *Engineering Psychology and Human Performance*, Prentice-Hall, New York, 1999.
23. A. Kuehberger, "The influence of framing on risky decisions: A meta-analysis," *Organizational Behavior and Human Decision Processes*, **75**, 23–55, 1998.
24. D. Kahneman and A. Tversky, "Prospect theory: An analysis of decision under risk," *Econometrica*, **47**, 263–291, 1979.
25. J.V. Schultz, "A framework for military decision making under risks," thesis, School of Advanced Airpower Studies, Maxwell AFB, AL, 1997.
26. J. Levy, "Applications of prospect theory to political science," working paper, Rutgers University, New Brunswick, NJ, 2002.
27. R.B Skov and S.J. Sherman, "Information-gathering processes: Diagnosticity, hypothesis-confirmatory strategies, and perceived hypothesis confirmation," *Journal of Experimental Social Psychology*, **22**, 93–121.
28. C.R. Mynatt, M.E. Doherty, and R.D. Tweney, "Consequences of confirmation and disconfirmation in a simulated research environment," *Quarterly Journal of Experimental Psychology*, **30**, 395–406, 1978.
29. N. Dixon, op. cit.
30. A. Wrangham, "Is military incompetence adaptive?" *Evolution and Human Behavior*, **23**, 3–17, 1999.
31. P.K. Davis, "Appendix G.3. Protecting weak and medium-strength states: Issues of deterrence, stability, and decision making," in National Research Council, *Post-Cold War Conflict Deterrence*, National Academy Press, Washington, 161, 1997.
32. W. Casscells, A. Schoenberger, and T. Grayboys, "Interpretation by physicians of clinical laboratory results," *New England Journal of Medicine*, **299**, 999–1000, 1978.
33. G. Gigerenzer, "Why the distinction between single-event probabilities and frequencies is important for psychology (and vice versa)," in *Subjective Probability*, G. Wright and P. Ayton, eds., 129–161, John Wiley & Sons, New York, 1994.
34. L. Cosmides and J. Tooby, op. cit.
35. J.St.B.T. Evans, et al., "Frequency versus probability formats in statistical word problems," *Cognition*, **77**, 197–213, 2000.

36. G.L. Brase, "Which statistical formats facilitate what decisions? The perception and influence of different statistical information formats," *Journal of Behavioral Decision Making,* **15**, 381–401, 2002.
37. R.J. Heuer, 1999, op. cit.
38. M.S. Cohen, "The naturalistic basis of decision biases," in *Decision Making in Action: Models and Methods*, G. Klein, et al., eds., 51–99, Ablex, Norwood, NJ, 1993.
39. N.J. Heacox, M.L. Quinn, R.T. Kelly, and J.W. Gwynne, "Decision support system for coalition operations: Final report," Technical Report 1886, SPAWAR Systems Center, San Diego, 2002.
40. A. Tversky and D. Kahneman, "Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment," *Psychological Review*, **90**, 293–315, 1983.
41. K. Fiedler, "The dependence of the conjunction fallacy on subtle linguistic factors," *Psychological Research*, **50**, 123–129, 1988.
42. A. Tversky and D. Kahneman, 1983, op. cit.
43. L.L. Lopes and G.C. Oden, "The rationality of intelligence," in *Probability and Rationality: Studies on L. Jonathan Cohen's Philosophy of Science*, E. Eells and T. Maruszewski, eds., 199–223, Rodopi, Amsterdam, 1991.
44. M. Bar-Hillel and E. Neter "How alike it is versus how likely it is: A disjunction fallacy in probability judgments," *Journal of Personality and Social Psychology*, **65**, 1119–1131, 1993.
45. K. Sawyer, "Witness chides NASA in assessing risk; ex-space official cites review of problem-plagued 1999 Columbia mission," *The Washington Post*, A08, March 7, 2003.
46. A. Tversky and D. Kahneman, 1974, op. cit.
47. R.M. Hogarth, "Beyond discrete biases: Functional and dysfunctional aspects of judgmental heuristics," *Psychological Bulletin*, **90**, 197–217, 1981.
48. J. Shanteau, "Cognitive heuristics and biases in behavioral auditing: Review, comments, and observations," *Accounting, Organizations, and Society*, **14**, 165–177, 1989.
49. R. Lipshitz, G. Klein, J. Orasanu, and E. Salas, "Taking stock of naturalistic decision making," *Journal of Behavioral Decision Making*, **14**, 331–352, 2001
50. M.S. Cohen, "The bottom line: Naturalistic decision aiding," in *Decision Making in Action: Models and Methods*, G. Klein, et al., eds., 265–269, Ablex, Norwood, NJ, 1993.
51. R.T. Kelly, J.G. Morrison, and S.G. Hutchins, "Impact of naturalistic decision support on tactical situation awareness," *Proceedings of the 40th Human Factors and Ergonomics Society Annual Meeting*, Philadelphia, 1996.
52 R. Samuels, S. Stich, and M. Bishop, "Ending the rationality wars: How to make disputes about human rationality disappear," in *Common Sense, Reasoning and Rationality*, R. Elio, ed., 236–268 , Oxford University Press, New York, 2002.
53. D. Serfaty, J. MacMillan, E. Entin, and E. Entin, "The decision-making expertise of battle commanders," in *Naturalistic Decision Making*, C.E. Zsambok and G. Klein, eds., 233–246, Lawrence Erlbaum, Mahwah, NJ, 1997.
54. J.G. Morrison, R.T. Kelly, R.A. Moore, and S.G. Hutchins, "Implications of decision making research for decision support and displays," in *Making Decisions Under Stress: Implications for Individual and Team Training*, J.A. Cannon-Bowers and E. Salas, eds., 375–406, APA Books, Washington, 2000.
55. Mark Mandeles, T.C. Hone, and S.S. Terry, *Managing "Command and Control" in the Persian Gulf War*, Praeger, Westport, CT, 1996.
56. B. Fischhoff, "Heuristics and biases in application," in *Heuristics and Biases: The Psychology of Intuitive Judgment*, T. Gilovich, D. Griffin, and D. Kahneman, eds., 730–748, Cambridge University Press, New York, 2002.
57. D. Serfaty, E.E. Entin, and R.R. Tenney, "Planning with uncertain and conflicting information," in *Science of Command and Control: Part II, Coping with Complexity*, S.E. Johnson and A.H. Levis, eds., 91–100, AFCEA Press, Fairfax, VA, 1989.
58. J.A. Sokolowski, "Modeling the decision process of a Joint Task Force Commander," doctoral thesis, Old Dominion University, Suffolk, VA, 2003.
59. A. Barsnick, "Modeling human elements in decision-making," master's thesis, Naval Postgraduate School, Monterey, 2002.

60. J.P. Kahan, D.R. Worley, and C. Stasz, *Understanding Commanders' Information Needs*, RAND, Santa Monica, 2000.
61. P. Ubel, "Identifying and reducing cognitive biases created by decision aids," www.med.umich.edu/pihcd/projects/decisionaid/decisionaid.htm, University of Michigan, 2002.
62. C.C. Cho, "Choose your weapons and targets," *The Edge*, **3**, 12, 1999.
63. M. Hura, et al., *Enhancing Dynamic Command and Control of Air Operations Against Time Critical Targets*, Santa Monica, RAND, 2002.
64. D. Pedersen, J.R. Van Zandt, A.L. Vogel, and M.R. Williamson, "Decision support system engineering for time critical targeting," *Command & Control Research & Technology Symposium*, Newport, RI, 1999.
65. Ibid.
66. D.A. MacGregor, "Future battle: The merging levels of war," *Parameters*, **22**, 33–47, 1992.
67. R. Lipschitz, "Knowing and practicing: Teaching behavioral sciences at the Israel Defense Forces Command and General Staff College, *Journal of Management Studies*, **20**, 121–140, 1983.
68. M.S. Cohen, et al., *Training Critical Thinking for The Battlefield. Volume I: Basis in Cognitive Theory and Research*, Cognitive Technologies, Arlington, VA, 2000.
69. B. Fischhoff, 2002, op. cit.
70 B. Fischhoff, "Prejudices about biases," working paper, Department of Social & Decision Sciences, Carnegie Mellon University, Pittsburgh, 1997.
71. J. Klayman and K. Brown, "Debias the environment instead of the judge: An alternative approach to reducing error in diagnostic (and other) judgment," *Cognition*, **49**, 97–122, 1993.
72. D. Arnott, "Decision biases and decision support systems development," Working Paper 2002/04, Decision Support Systems Laboratory, Monash University, Melbourne, 2002.
73. US Air Force, Deputy Chief of Staff, Plans and Operations, *JFACC Primer*, 2nd Edition, Washington, 1994.
74. D. Arnott, 1998, op. cit.
75. L.J Sanna, N. Schwarz, and S.L. Stocker, "When debiasing backfires: Accessible content and accessibility experiences in debiasing hindsight," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **28**, 497–502, 2002.