AD-A170 952 A NEW METHODOLOGY FOR MODELING NATIONAL COMMAND LEVEL DECISIONMAKING IN W..(U) RAND CORP SANTA MONICA CA P K DAVIS ET AL. JUL 86 RAND/R-3290-NA MDA903-85-C-0030 1/2

UNCLASSIFIED F/G 5/10 NL

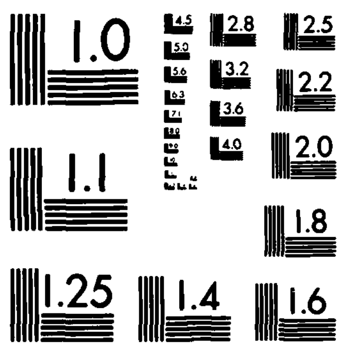

!CROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

(12)

AD-A170 952

A New Methodology for Modeling National Command Level Decisionmaking in War Games and Simulations

Paul K. Davis, Steven C. Bankes, James P. Kahan

A Report from
The Rand Strategy Assessment Center

Rand

The research described in this report was sponsored by the Director of Net Assessment, Office of the Secretary of Defense, under Contract No. MDA903-85-C-0030.

Library of Congress Cataloging in Publication Data

Davis, Paul K., 1943-
A new methodology for modeling national command level decisionmaking in war games and simulations.

"July 1986."
"R-3290-NA."
Bibliography: p.
1. War games—Decision making—Mathematical models.
2. War games—Decision making—Simulation methods.
1. Bankes, Steven C. II. Kahan, James P., 1942-
III. United States. Dept. of Defense. Director of
Net Assessment. IV. Rand Corporation. V. Title.
U310.D384 1986 355.4'8'0151 86-13863
ISBN 0-8330-0735-1

Published by The Rand Corporation

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER R-3290-NA	2. GOVT ACCESSION NO. AD-A170952	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) A New Methodology for Modeling National Command Level Decisionmaking in War Games and Simulations		5. TYPE OF REPORT & PERIOD COVERED Interim
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) P. Davis, S. Bankes, J. Kahan		8. CONTRACT OR GRANT NUMBER(s) MDA903-85-C-0030
9. PERFORMING ORGANIZATION NAME AND ADDRESS The Rand Corporation 1700 Main Street Santa Monica, CA 90406		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Director of Net Assessment Office of the Secretary of Defense Washington, DC 20301		12. REPORT DATE July 1986
		13. NUMBER OF PAGES 117
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for Public Release; Distribution Unlimited

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

No Restrictions

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

War Games
Simulation
Mathematical Models
Decisionmaking

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

see reverse side

DD FORM 1 JAN 73 1473

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

This report describes and illustrates a methodology for modeling National Command Level (NCL) decisionmaking in large-scale crisis and conflict involving both superpowers--a methodology that progresses from abstract concepts about superpower objectives and strategy through the step-by-step procedures for building an operational computer program. Its two key components are defining an "image" of Soviet or U.S. decisionmaking, and moving from that imprecise image to a precise and coherent computer program. The report's sections outline the modeling approach; review the most important concepts underlying the approach; describe the systematic definition of alternative coherent images of the Soviet or U.S. NCL; describe building an operational computer program that is both transparent and able to explain its own decisions; and discuss initial experiences using prototype versions of the computer models.

R-3290-NA

A New Methodology for Modeling National Command Level Decisionmaking in War Games and Simulations

Paul K. Davis, Steven C. Bankes,
James P. Kahan

July 1986

Prepared for the
Director of Net Assessment,
Office of the Secretary of Defense

A Report from
The Rand Strategy Assessment Center

Rand
1700 MAIN STREET
P.O. BOX 2138
SANTA MONICA, CA 90406-2138

PREFACE

The work described in this report was accomplished as part of a larger research program in the Rand Strategy Assessment Center (RSAC), primarily under the sponsorship of the Director of Net Assessment in the Office of the Secretary of Defense. The report can be read without prior knowledge of the RSAC's research and is intended to be of interest to a diverse audience, including those interested in military strategy, decisionmaking theory, simulation, and artificial intelligence. Parts of the report are, however, inherently technical, and some readers may wish merely to skim them. Finally, because the work breaks new ground in several domains where there is little published work, the authors would be especially interested in readers' comments and suggestions, which should be directed to the RSAC's director, Paul K. Davis.

SUMMARY

OBJECTIVES

This report describes and illustrates a methodology for modeling National Command Level (NCL) decisionmaking in large-scale crisis and conflict involving the superpowers. We have developed prototype NCL models using this methodology and are now developing second-generation models that have significant potential for a number of diverse activities:

- *Evaluating military strategies, forces, and command-control systems* using game-structured political-military simulations;
- *Providing decision aids or stand-ins* for human players in political-military war games being conducted for training or to explore new strategic concepts; and
- *Studying, with some degree of rigor, alternative views of deterrence, escalation control, and war termination.*

In the first of these, the NCL models will be one part of a much larger game-structured simulation system being developed by the Rand Strategy Assessment Center as an improved tool of global, integrated, strategic analysis. In this role, the NCL models can be thought of as generators of plausible scenarios in which military forces and strategies can be evaluated—scenarios that may begin with crisis and extend through general nuclear war, or that may begin with crisis and terminate without cataclysmic warfare (but with the participants constantly aware of the nuclear shadow).

As decision aids in human games, the NCL models can be used for *situation assessment* and to help structure the discussion of issues, options, and rationale for possible decisions. NCL models will also be used to stand in for Red teams—i.e., to simulate Red decisionmaking while human teams develop Blue decisions (or vice versa). This will allow greater control over the threats presented to Blue teams and will allow Blue teams to be exposed to "standard" problems and various types of Soviet behavior thought significant by Soviet specialists.

The third activity focuses more on the *building* of models than on their operation. By contrast with the scientific disciplines, the study of international behavior in crisis and conflict lacks a highly developed set of analytic techniques by which to tighten arguments and communicate them to others unambiguously. The mechanism of building and

discussing NCL models provides new opportunities in this regard. Building such models forces one to specify the *context* of assumptions about international behavior in substantial detail, and this often clarifies or resolves disagreements ("Well, perhaps they *would* use nuclear weapons in that situation, but you and I obviously disagree about how plausible that situation is."). Furthermore, because uncertainties about future national behavior can be reflected with *alternative* NCL models, it is possible to characterize a range of plausible behavior patterns without requiring convergence on a single best estimate.

It should be possible, in particular, to use the building of NCL models as a mechanism for studying deterrence, escalation control, and war termination with some degree of rigor. Past studies of such matters have been unduly dominated by qualitative essays on the one hand, and on the other hand by quantitative approaches focused on oversimplified statements of the problem. By building and discussing NCL models concerned with both qualitative issues (e.g., alliance cohesiveness or the quality of strategic warning) *and* quantitative issues (e.g, the vulnerability of one's strategic nuclear forces), it should be possible to look more deeply and rigorously into the issues. It should also be possible to give greater prominence to such issues as asymmetries in the perspectives of different nations, uncertainties stemming from imperfect intelligence and imperfect understanding of the other participants in crisis and conflict, and the values held by real-world leaders.

SUMMARY OF METHODOLOGY

There are two aspects of the methodology for modeling NCL decisionmaking: developing a reasonably clear image of the Soviet Union or the United States, and translating that image into a computer program. As mentioned above there are fundamental uncertainties about Soviet and U.S. behavior, which causes us to build alternative models referred to as alternative "Ivans" and "Sams." However, we build one model at a time to assure some degree of coherence. The first issue, then, is characterizing a particular Ivan and Sam.

Characterizing Ivans and Sams

Several procedures assist the analyst in developing a strong image of a particular Ivan or Sam before attempting to develop a detailed model. The first step is to write a short essay describing the model's intended world view, grand strategy, value system, and temperament. The second step is to fill out a formal checklist of attributes that deal with such matters as willingness to use nuclear weapons and perception of the opponent. After this exercise, the analyst may wish to revise the essay, but it should be kept short. A third step is to fill out a matrix showing possible changes in the conflict state to rank order those escalations and deescalations the model should be made to treat early in its development.

The next step attempts to draw out notions of grand strategy in conflict. The analyst considers a number of key situations (e.g., one in which the Red NCL has already decided to invade Europe and is now contemplating his grand strategy for doing so) and sketches out decision trees showing, roughly, what his Ivan or Sam would be thinking about in such a situation (e.g., possible branches in the conflict for which he must be prepared). The purpose here is not to be complete but rather to sharpen the basic image. At this point, then, the analyst should be able to jot down a number of guiding principles that appear to have been used in the earlier portions of this Ivan- or Sam-defining exercise. These principles are first approximations of rules dictating when the model should or should not be willing to escalate and terminate, what its objectives would tend to be, and how it might try to carry out those objectives.

Since this part of the methodology is to stimulate *first-order* thinking, it should be kept brief and accomplished quickly because the high payoff comes early and attempts to push it too far will prove frustrating. The measure of success is in the next step: As the analyst begins actually to build or adapt NCL models, does he constantly have to rethink his concept of the Ivan or Sam, and is he able to maintain coherence?

Building a Model of Ivan or Sam

In our methodology, an NCL model consists of a large number of discrete *rules* defining how, for example, the particular Ivan or Sam would characterize a particular situation, how he might characterize his opponent given the opponent's behavior so far in the conflict, and how he would decide on such issues as escalation. The rules take the form of If <condition> Then <decision, judgment, or action>. For example, a simple rule might be "If NATO-is-mobilizing Then Let Strategic-warning be Eur-gen-conv", meaning that the particular Ivan considers NATO mobilization to be strategic warning of a conventional war in Europe. Without such explicit rules, the model can have no concept of "strategic warning" or any of the other issues central to strategic thinking in crisis and conflict. Very large numbers of such rules are needed if the model is to be at all intelligent.

The problem in developing a methodology, then, is to define a conceptual architecture to guide rule-writing and within which to organize the rules once written. Without such an architecture it would be altogether impractical to build a complex rule-based NCL model: Where would one begin, when would one be done, and how would one know?

In developing a conceptual architecture, we were strongly influenced by the objectives for which the models were being developed—objectives that require us to *understand* the models' logic and that define the models as mechanisms for studying complex issues amidst uncertainty. Some of our principal requirements were:

- *Realism*: The reasoning exhibited by NCL models should be natural in human terms, with the models focusing on the same type of variables that real-world decisionmakers would consider important in the crises and conflicts being simulated.
- *Transparency*: A model's logic should be understandable and human-like in terms of its decision process and individual judgments.
- *Flexibility*: The first model developed should really be a framework for a diversity of models representing alternative Ivans and Sams; it should be possible to reflect diversity in grand strategies, value systems, perceptions, and sheer competence.
- *Evolutionary potential*: Though a first-generation NCL model would surely be highly simplified, it should provide a good baseline for more sophisticated models.
- *Ease of use*: It should be possible to review and adapt NCL models without being a proficient computer programmer.

With this prelude, then, **the model architecture can be summarized as follows in terms of the way we organize rule-writing**. We organize:

- *By the steps in an understandable decision process*: As shown in Fig. S.1, the NCL models choose a course of action called an analytic war plan, along with certain details of that plan, by proceeding through a particular reasoning process that begins with situation assessment. The rules for each step of that process are grouped together.
- *By current state of the conflict*: Within each step of the process indicated in Fig. S.1, the rules are organized by distinguishing among different levels of global conflict and, within that structure reminiscent of an escalation ladder, by distinguishing more

finely among conflict states in the individual theaters worldwide.

- *By a hierarchy of variables*: Within the group of rules applying to a particular step of the decision process for a particular class of world situation, an NCL model attempts to reason in a strategic manner by using aggregated concepts such as status, prospects, and risks. The value of each such variable may depend on the values of several lower-level variables, which in turn may depend on the values of still lower-level variables. Ultimately, the NCL's assessment of something like risks may, for example, depend on whether the opponent is preparing to escalate and whether his own nation's intelligence and communication systems can detect and interpret those preparations. In addition, the assessment of risks will depend on judgments about the opponent and many other factors as well.

The character of the NCL's intended decisionmaking is something like that of the following hypothetical decisionmaker who, after having been deluged with information in the form of briefings, memoranda, and personal advisories, sums up as follows:

> Well, gentlemen, if I understand what you have been telling me these last few hours, and if I try to patch together some of the pieces that came out one at a time in our meeting, then it seems that our *current situation is pretty good*—we have achieved our principal objectives, although not everything we had hoped for. We could push on, but *prospects for further progress appear only marginal* and *there appear to be big risks*. If that is all correct, then I conclude we should begin to consolidate our gains and wind down our actions. Do you agree?

To summarize, then, **the writing of rules and the organization of rules once written are based on a process model of decision, on a characterization of alternative states of the world, and on the use of variable hierarchies.** To the maximum extent possible, one develops variants of baseline NCL models by merely changing particular rules within the overall structure or by adding complexity to the rule hierarchies. However, when necessary it is possible to change the representation of world states, to vary the manner in which alternative decisions are compared, and so on. Thus, the model framework has substantial flexibility.

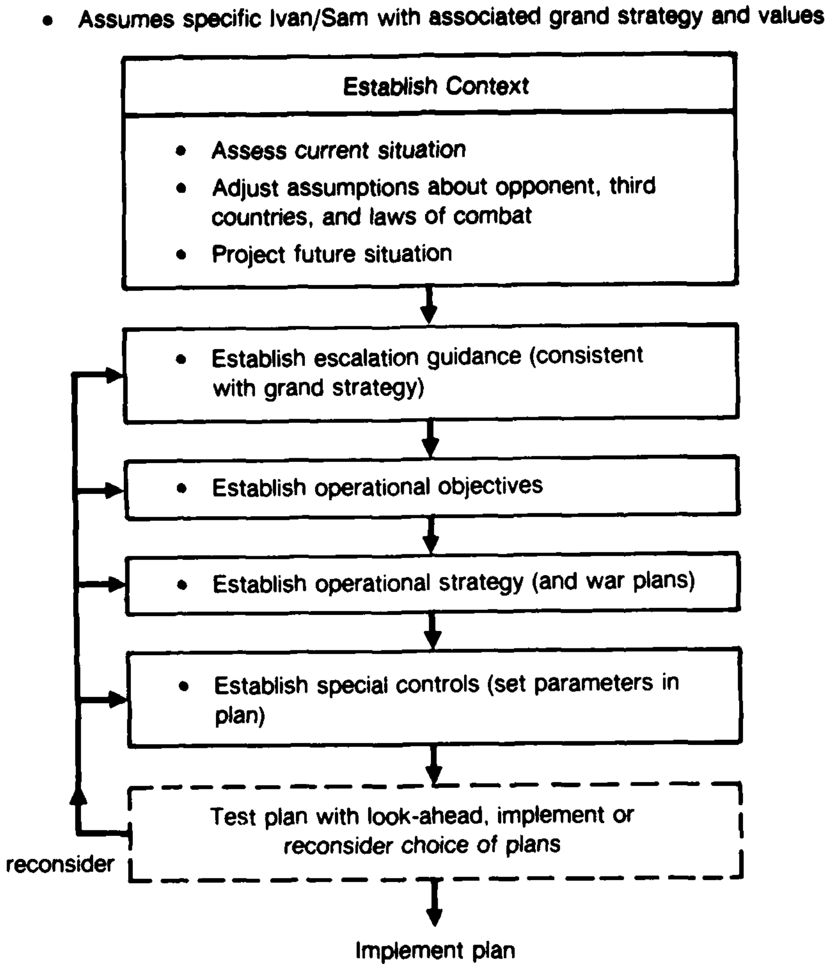

Fig. S.1—A process model of NCL decisionmaking

Observations About the Approach

The approach defined here has several important features, as suggested by Fig. S.1 and as elaborated in the text. These include:

- A top-down, global, strategic perspective.
- A human-like logic dependent more on heuristic decision processes than on the more formal techniques of quantitative decision analysis, but with many features of the so-called ideal decisionmaker of cognitive science (e.g., examination of multiple options, search for tie-breaking information, and follow-up with feedback).
- "Look-ahead" projections, using an NCL model's *perceptions* of reality, as part of the decisionmaking process.
- Treatment of perceptions based on current assumptions about the nature of the opponent, the status of conflict, the positions of third countries, and the laws of war—assumptions that can be wrong because of intelligence or command-control failures, biases, or simple misjudgments.
- Situation assessment based on a top-down hierarchical treatment of information.

This approach is by no means uniquely correct in the sense of simulating the thought process of actual decisionmakers. However, we claim that the framework is reasonably natural, logical, understandable, and flexible enough to accommodate a broad range of realistic Soviet and U.S. behavior patterns. Thus, although the framework will be enriched with the benefits of experience and ongoing research, we believe the current framework is more than adequate to shift the limiting factor in the development of NCL models from matters of technique to matters of substance.

NCL Models As Artificial Intelligence

Technically, the research described here represents a unique and ambitious application to policy analysis of artificial intelligence techniques, primarily those associated with so-called knowledge-based systems. Consistent with generic goals of such work, our models are flexible, transparent, and comprehensible as a whole—with the models providing automatic first-order explanations of their decisions. Analysts with only modest knowledge of programming can look directly into the computer code, read and understand the important decision rules, and make significant changes themselves. This is possible by virtue of a new fast programming language called Rand-Abel™, which allows rules to be written in English-like expressions or in decision tables identical with those an analyst might develop on his worksheet.

In many respects, the NCL models look like so-called *expert systems*. There are several important differences, however. At a technical level, our approach dictates the model's inferencing logic rather than relying upon general "search" techniques. Another difference is that there exist no real experts in the domain we are working in—i.e., there are no people to whom we can go for a reliable and reasonably complete set of rules that have been proven valid empirically in a diversity of superpower crises and conflicts (there have been *very* few superpower crises and no superpower conflicts). To the contrary, much of our work is inherently analytic rather than empirical.

STATUS AND PROSPECTS FOR FUTURE WORK

As mentioned above, prototype NCL models exist and operate. Although the prototypes have not been exposed to thorough scrutiny, we have described them in general terms to a large number of people—primarily through briefings and on-line demonstrations. And, since the prototypes were completed, we have conducted several "seminar war games" focused on NCL issues. These consisted of small groups talking through what the issues and decisions might be in a variety of high-level crisis and conflict. Some participants had many years of experience with political-military war gaming and/or policymaking. Although our conclusions on this are inherently and unabashedly subjective, we are now convinced that the prototype models go far toward capturing the issues those participants believe are important.

We are now developing second-generation models. The emphasis in this work is on enriching the models' substantive content, because the models are already relatively strong technically (although many interesting extensions are possible). Currently, we are aware of the need to enrich the models in at least the following respects:

- Enhance sensitivity of decision rules to perceptions of the opponent, third countries, and laws of war.
- Enhance sensitivity to the time dimension (e.g., the pace of events and its relationship to rationality).
- Enhance sensitivity to command and control effects generally, which will be increasingly feasible as other models dealing with command and control emerge from parallel research on the full game-structured simulation.

- Broaden the class of scenarios and detailed situations for which the rule-sets are potentially applicable (e.g., conflicts focused on the Far East or Middle East).
- Develop a reasonable *set* of Ivans and Sams representing different grand strategies and temperaments. Explore in some depth the issue of U.S.-Soviet asymmetries.
- Improve the quality of decisionmaking by performing multiple look-aheads (e.g., testing a plan with both best-estimate and worst-case assumptions about the opponent's likely behavior).
- Improve the quality of decisionmaking by making explicit plan comparisons on the basis of look-aheads or more limited projections.
- Add optional stochastic features to the most critical of decision points.

As work in these areas proceeds, it will be possible to perform more rigorous experiments comparing model decisions with those of humans and observing the degree to which human teams will make different decisions if prompted by suggestions from models accompanied by logical explanations.

Finally, we should add the comment that much of our work could have analogues in other domains of policy analysis. Almost all large corporations have strategic planning functions, as do government agencies. Our work represents an unusual and probably unique effort to com')ine in a large-scale complex problem area the techniques of both rule-based heuristic modeling and traditional time-oriented simulation of processes and events, and to do so in a game-structured paradigm. It seems likely that similar efforts would prove useful in other domains in which one sees adversarial processes (or dynamic response of the "environment"), and a combination of organizational and strategic behaviors. We hope to explore some of these issues in future work.

ACKNOWLEDGMENTS

The authors appreciate the thoughtful review comments of Paul Bracken, Stephen Cimbala, Alan Platt, Peter Stan, and James Tritten. Both Peter Stan and Randall Steeb contributed to the initial work on National Command Level modeling, and Robert Weissler has greatly improved the prototype computer programs for the NCL models.

CONTENTS

FIGURES

I. INTRODUCTION

OVERVIEW

This report describes and illustrates a methodology for developing alternative models of National Command Level (NCL) decisionmaking in large-scale crises and conflicts involving both superpowers—a methodology that leads the reader from abstract concepts about superpower objectives and strategy through the step-by-step procedures for building an operational computer program.

The purpose of this overview is to touch upon the following questions: (1) Why would one want to develop NCL models? (2) Is it feasible to do so? and (3) How would one measure the quality of such models once developed? Having discussed those matters, we shall then provide the reader with more background on the origins of the effort and its relationship to other work by the Rand Strategy Assessment Center (RSAC).

Motivation for Developing NCL Models

It is natural to ask why one would need or want NCL models, since strategic analysis has been conducted for years without such constructs and it is evident that the NCL models deal with matters about which there is enormous uncertainty—thus casting doubt on the models' usefulness. Why, then, do we bother? The principal reasons are as follows.

Analysis using simulations of crisis and conflict. The principal motivation for developing NCL models is the desire to use game-structured political-military simulation as a means for assessing alternative military forces and strategies. The NCL models are part of a much larger simulation system that will be described below. Within that system, the role of the NCL models is, in a sense, to make the top-level decisions that generate plausible scenarios for which the evaluations can be conducted—scenarios extending all the way from crisis through general nuclear war in some instances. Traditionally, strategic analysis has been highly compartmentalized, with some studies dealing with lower-level crises, others with conventional conflicts, others with theater-nuclear conflicts, and still others with intercontinental war. Although there are distinct advantages to that approach, there is clearly a need for integrative work as well. The political-military simulations should make such work possible (Davis and Winnefeld, 1983).

In addition to using simulations to evaluate forces and strategies, it should be possible with NCL models imbedded in a larger simulation to examine aspects of strategic command-control that are usually ignored altogether. That is, although there exists considerable work on technical aspects of command and control, particularly on communications, tactical warning, and attack assessment, it has been notoriously difficult to deal analytically with those aspects of strategic command-control involving top-level human decisions. Some of the best insights have come from human war games and it seems reasonable to believe that analytically controlled war games or game-structured simulations should also pay dividends in this area (Davis, Stan, and Bennett, 1983).

Gaming and interactive simulation for training and exploration. A second motivation for NCL models is the desire to provide players in human war games with decision aids and, sometimes, to replace entire teams with reasonable decision models. Gaming has long been a major technique for exploring strategic concepts and for de facto training. However, human teams vary in their sophistication and discussion within human teams can be chaotic. NCL models should provide structure for discussion, *propose* possible decisions, and provide rationale. The human teams may then learn from the experience of their predecessors, whose wisdom could be reflected in the models, and see an organized portrayal of information. They may choose to ignore the models, or to modify the suggested decisions, but the models should nonetheless be useful as aids.

In other instances, it should be useful to replace an entire team (typically, a Red team) with an NCL model so that the war game can be better controlled and more expeditiously conducted. One of the chronic difficulties in war gaming has been the paucity of Red specialists available to play the Red leadership. Furthermore, if the purpose of a war game is to expose the Blue team to particular problems and types of behavior, then one wants to control the Red team's behavior—i.e., one really wants a reproducible model.

The rigorous study of deterrence, escalation control, and war termination. There are clear reasons, then, to desire NCL models for their potential value in games and simulations. Beyond those, it seems that the very process of *developing* NCL models should be valuable as a mechanism for inserting rigor into a subject area dominated by qualitative essays at one extreme and by overly quantified studies at the other (Davis and Stan, 1984). It can be argued (Davis, 1986) that there have been no major *analytic* advances in the study of deterrence, escalation control, and war termination since the seminal work of the 1950s and early 1960s, and that the process of building NCL models could provide the structure and rigor needed to tighten arguments and clarify points of disagreement. Especially significant is the need in NCL models to specify the *context* of decisions—something often left only partially defined in debates and essays and something at the heart of many disagreements.

It is evident from considering all of these reasons for having NCL models that it is necessary to have a diversity of models. That is, it is necessary to represent alternative plausible images of Soviet and U.S. decisionmaking (using models we refer to as alternative "Ivans" and "Sams," respectively). It is not necessary that specialists agree on a particular behavior pattern; instead, we can have alternatives, which also has advantages for improving communications among analysts.

Feasibility and Appropriateness of NCL Models

There are several reasons for doubting a priori that one can build useful NCL models. For example, it could be argued that the behavior of a nation in crisis and conflict would consist of many decisions, any of which could go in several directions—i.e., the model would have to consist of a great many decision rules, each of which would have a highly random component. An article of faith in our work is that so long as a nation's leadership is reasonably stable and rational, there will be *patterns* of behavior correlating the individual decisions. Thus, we need not consider all possible combinations of decisions, but only selected patterns (and some excursions to reflect uncertainty about particularly difficult decisions).

A second basis for doubt about NCL modeling is the concern that perhaps we should be concerned less with modeling *likely* behavior than with analyzing *capabilities* assuming a malign opponent trying to do his worst. This translates into the classic argument that strategic analysis should focus on capabilities rather than intentions. Here, our attitude is that previous analysis has focused on the extreme of looking *only* at capabilities, without adequate regard for intentions, doctrines, mindsets, and perceptions. For many purposes, capability-oriented analysis is altogether appropriate, but for others it is not. In real-world crises, for example, it is *essential* to know our opponent, and even to affect his image of us.[1] A competitor who always looks to the worst case can readily find himself paralyzed. Of course, it is also true

[1]The effort to affect an opponent's image of oneself is part of what is sometimes called *reflexive control*, and leads to strategies very different from those arising in usual game-theory studies. See, for example, Lefebvre and Lefebvre (1984, 1985), which describe work begun by the Lefebvres in the Soviet Union during the 1960s. More generally, reflexive control involves affecting the opponent's decision process by affecting his image of reality.

that by underestimating the enemy one can suffer grievously (as, for example, Russia suffered when Stalin's assessment of Hitler's intentions proved faulty in the weeks immediately preceding the German invasion of Russia). The point, then, is that real-world decisionmakers *must* operate with an image of likely opponent behavior, and failure to model that behavior is tantamount to accepting an implicit model that may well be wrong.[2]

Criteria for Evaluating NCL Models

Given motivation to build NCL models and the conclusion that doing so might be feasible, it still remains to define more precisely what the models should be expected to accomplish. It is evident, for example, that the enormous uncertainties associated with national decisionmaking preclude building NCL models that could be "validated" in the sense that term is used when testing some computer program for solving differential equations.

In this regard, the first point to make is that it is by no means necessary that the models *always* be sensible—they are to be used in interactive man-machine environments in which it is quite reasonable to expect the human analyst to override (and find humor in) some of the model's pronouncements.[3] A more appropriate measure of success is whether the models *usually* perform as well as human control teams in war games conducted for research, training, or analysis. A second measure is whether the variables examined by the models are the variables considered important in large-scale crises and conflicts by people with high-level experience. A third measure is whether an analyst comparing his judgments with those of the models under a variety of circumstances finds that the models help him to think more sharply and to understand better the key issues that separate his views from those of others. Although major improvements will be necessary, preliminary indications are that even the relatively simple first-generation models have merit by all three criteria so long as they are used for their intended purposes, which are essentially analytic.

This, then, constituted the motivation for developing NCL models, our basis for believing that the effort could be successful, and our criteria for judging success. Let us next discuss the context in which the work has been performed.

---

[2]This is similar to the admonition of Forrester regarding the importance of including subjective estimates of "soft" variables in policy analysis, an admonition that was (and to some extent remains) surprisingly controversial. See, for example, Forrester (1969) and Randers (1980), which discuss Forrester's System Dynamics approach to policy analysis in both philosophical and technical terms. Davis (1985a) discusses the implications of this attitude about soft variables for assessments of the NATO-Warsaw Pact balance.

[3]In this and several other respects the NCL models have much in common with *expert systems*, which it is now recognized should be regarded more as intelligent assistants than as replacements for humans.

BACKGROUND

The research described here is part of a much larger effort by the RSAC to develop new methods of strategic analysis combining features of war gaming and analytic modeling (Davis and Winnefeld, 1983). A central element of that effort is a game-structured simulation in which all of the human teams of a traditional political-military war game can be replaced by models as indicated in Fig. 1.1.[4] Although the prototype system is incomplete and simplified in many respects, the limiting factors are now subject-area specific rather than technical—i.e., at a technical level, the system is already a reality.

The function of the Red and Blue Agents shown in Fig. 1.1 (i.e., the models representing the Soviet Union and United States, respectively) is to choose, implement, and adapt appropriate *analytic war plans*. These plans specify military and political actions simulated by the Force Agent, which keeps track of worldwide forces, computes the

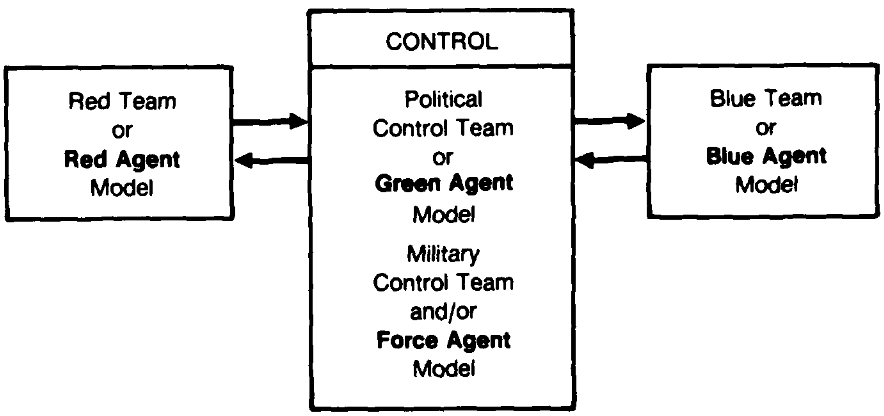

Red Agent: Makes decisions for Soviet Union and Warsaw Pact
Blue Agent: Makes decisions for United States and NATO
Green Agent: Makes decisions for nonsuperpower countries
Force Agent: Simulates results of combat and other military operations; also, controls simulation time

Fig. 1.1—Game-structured simulation

[4]Earlier RSAC publications use Scenario Agent for what is now referred to as the Green Agent.

results of military operations including combat, and keeps the overall simulation's clock. All of these actions take place in the context of third-country decisions simulated by the Green Agent (Schwabe and Jamison, 1982; Shlapak, Schwabe, Lorell, and Ben-Horin, 1985).

The simulation is highly interactive and most research is conducted with humans in at least one position (e.g., a Blue team or analyst testing strategies against an automated Red Agent). The models codify the insights gained from human play, increase the efficiency of such play, and are essential for conducting controlled analysis—especially for complex simulations involving multiple forces and theaters.

As shown in Fig. 1.2, the Red and Blue Agents consist of submodels representing the NCL and a hierarchy of military commands. For Blue, the Global Command Level (GCL) model has functions akin to those of the Joint Chiefs, with some wartime functions of the White House staff and State Department as well; the Supertheater Command Level (SCL) (of which there may be none or several) might correspond, for example, to the NATO Supreme Allied Commander in Europe (SACEUR) and the U.S. CINCEUR; the Area Command Levels (ACLs) correspond to commanders for the various individual theaters, such as central, northern, and southern Europe, Southwest Asia, the Pacific, and Atlantic. We consider space and the intercontinental arenas to be separate theaters as well.[5]

The NCL model must determine the overall analytic war plan, which amounts to specifying national military strategy and the numerous constraints on military and political actions. The SCL and ACL models fill in details of the analytic war plan, and make decisions about how to adapt the plan to circumstances and how to manage forces on a continuing basis. So long as the broad features of the NCL-specified plan continue to appear valid, it is the military command levels that make the decisions as the simulation proceeds. If, however, events in the simulation run counter to basic precepts of the plan (e.g., if the opponent escalates the level of conflict), then the NCL model must consider changing or modifying it. Also, under some circumstances the military level models request permission within a given analytic war plan to exert greater independence—i.e., they ask for some degree of increased delegation. Again, the NCL model must make the decisions.

[5]In applications work we use structures closer to real-world commands. Figure 1.2 applies only to prototype models. Note also that some ACL commanders have the same status as some supertheater commanders (e.g., CINCPAC has the same status as CINCEUR, even though the latter has subordinate theater commanders).

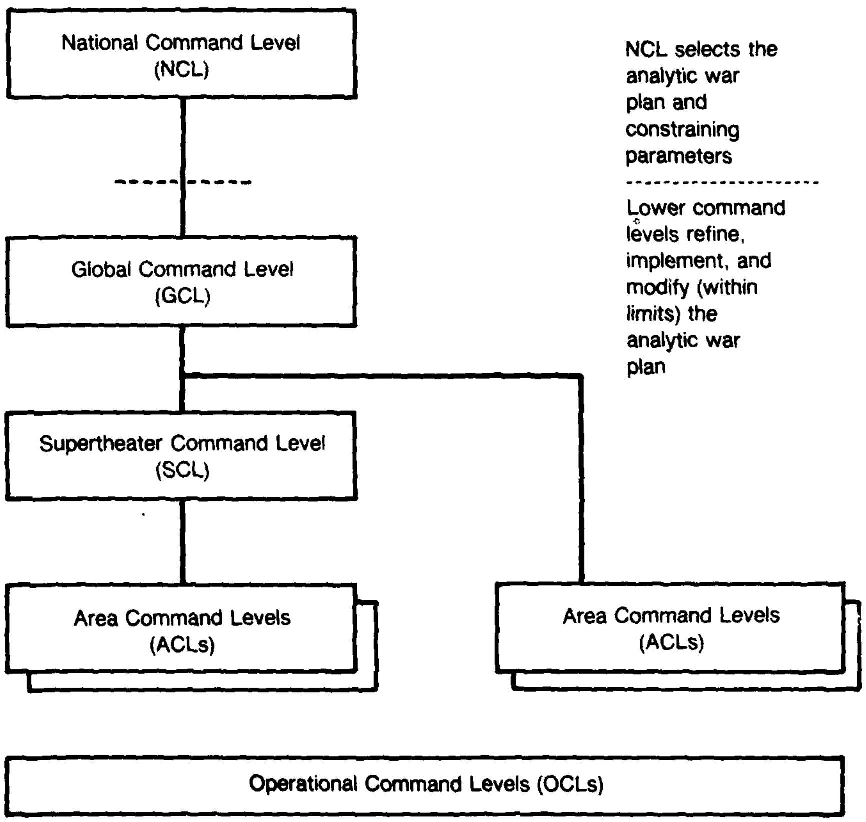

Fig. 1.2—Simplified view of command levels within the Red and Blue Agents (Red may have no SCLs; Blue may have several)

LIMITATIONS OF SCOPE

It is important to recognize explicitly the limitations of scope in any analytic effort. Most RSAC work to date, and this report in particular, deals only with large-scale crises and conflicts between the superpowers, especially those situations that might escalate to nuclear war. It therefore does not concern itself with peacetime decisionmaking, nor even with decisionmaking in the types of crises that have actually occurred since 1945, with the sole exception of the Cuban Missile Crisis.[6] The models treated here focus on fundamental political-military considerations and national values. They do not consider explicitly the powerful—even dominating—influences that short-term and highly personal political pressures have on real-world leaders during more ordinary crises and wars. Although the techniques used in this report could be extended to include such political factors, it would have been inappropriate to do so for the purposes that motivated the current research—e.g., evaluating in peacetime alternative force structures and military strategies.

THE REPORT IN OUTLINE

With this background, let us now summarize briefly what follows in the remainder of the report. It is important to recognize that the methodology defined here has two components: defining an "image" of Soviet or U.S. decisionmaking (i.e., a mental model of a particular Ivan or Sam), and moving from that imprecise image to a precise and coherent computer program. Indeed, it is useful to distinguish an intermediate step, which is to define the model to be implemented as a computer program.[7] We shall begin, in fact, by discussing the modeling approach. Section II discusses the most important concepts underlying our approach, including alternative concepts of the decisionmaking process and the general analytic framework we have adopted. Section III then describes how we systematically define alternative coherent images of the Soviet or U.S. NCL. Section III also illustrates the techniques we use to characterize alternative states of world conflict at a high level of aggregation—techniques that are used extensively in building the NCL models. Sections IV and V are somewhat different in character: They deal with how, for a given image of the Soviet Union or United States, we use the general modeling framework to build an operational computer program that is both transparent and able to explain its own decisions. These sections are more technical in nature and may be skimmed in a first reading. However, it is in these sections that one sees most concretely what the methodology entails. Finally, Sec. VI discusses our initial experience using prototype

---

[6]Although the Cuban Missile Crisis has been grist for the mills of numerous writers concerned with nuclear confrontations, it is notable that the Soviet Union never placed its forces on alert during that crisis, thereby suggesting, in retrospect, that the sides were never as close to the brink as is sometimes posited.

[7]The distinctions between model and program are well discussed in, for example, Zeigler (1984a).

versions of the computer models and indicates briefly some of the improvements currently under development. There is some deliberate redundancy across sections so that nontechnical readers can gain a coherent overview without reading Secs. IV and V and so that those sections can be reasonably self-contained.

II. BASIC CONCEPTS FOR DEVELOPING NCL MODELS

DISTINGUISHABLE ISSUES

There are many considerations in designing NCL models, most of them substantive rather than purely technical. The most important, in our experience, are:

1. Establishing the generic *character* of the models—i.e., decision style;
2. Developing a strategic framework within which to organize individual decision rules for *situation assessment*;
3. Characterizing particular *alternative NCL models* (alternative "Ivans" and "Sams"); and
4. Implementing the Ivans and Sams as *computer programs*.

We shall consider the first two of these in this section. The third item is discussed in Sec. III, and the last one in Secs. IV and V.

DEFINING THE GENERIC CHARACTER OF NCL MODELS

Alternative Approaches

As discussed in Sec. I, the purpose of the NCL models within the RSAC simulation is to choose analytic war plans and set certain parameters within them (e.g., parameters specifying rules of engagement or delegation of authority). This means, roughly, to pick a "strategy."

The problem, then, is to build a strategy-selecting model. There are numerous approaches about how best to perform decisionmaking and numerous theories about how real-world decisionmakers actually think. Some of the approaches to this problem, along with important references, include:

1. *Decision-analytic models* that specify formal methods for constructing utility functions and probability distributions for outcomes, which can then be used with algorithmic techniques to generate, within the formalized set of objectives and values, an "optimal" decision (Raiffa, 1970; Keeney and Raiffa, 1976);

2. *Game-theoretic models* that extend decision-analytic models to situations involving two or more interacting parties (Von Neumann and Morgenstern, 1953; Luce and Raiffa, 1957; Owen, 1982; Shubik, 1982; Kahan and Rapoport, 1984; and Brams, 1985). Game-theoretic strategies in two-sided conflicts consider explicitly the opponent's capability to "do his worst";
3. *Bureaucratic-politics models* that attempt to replicate the *process* of decisionmaking in the presence of competing agencies of government. Decisions result from compromise, coalition politics, and political power (Allison, 1971);
4. *Organizational process or so-called cybernetic models* that describe decisions of large organizations as relatively local adaptations to achieve predetermined goals. The adaptations may or may not be appropriate or optimal from a more global perspective given the new circumstances, but the organization does not attempt to make a new global evaluation (Simon, 1980, 1982; Steinbruner, 1974; Janis and Mann, 1977; Allison, 1971);
5. *Rule-based heuristic models* consisting of a compilation of plausible "If . . . then . . ." rules prescribing actions for diverse situations. Heuristic rules are usually considered to be rules that are good enough most of the time to guide actions. They are not guaranteed to be optimal. One example of such a model, Rand's Green Agent (Schwabe and Jamison, 1982; Shlapak, Schwabe, Lorell, and Ben-Horin, 1985), makes decisions for nonsuperpower nations in Rand's simulations of global conflict. See also the expert system literature (e.g., Hayes-Roth, Waterman, and Lenat, 1983); and
6. *Cognitive models* that focus on how decisionmakers receive and process information and how they cope with the problem of *bounded rationality* (Kahneman, Slovic, and Tversky, 1982, and references therein. See also the references above for organizational models).

The above approaches are by no means independent and rarely exist in pure form. For example, a cognitive model may derive its utilities from bureaucratic considerations or from heuristics. And, similarly, a model making heavy use of heuristics may also employ decision-analytic techniques to trade off conflicting values. There is also considerable overlap between those favoring organizational models and those favoring emphasis on heuristics rather than optimizing models.

Developing Requirements for RSAC NCL Models

With this abundance of concepts and techniques to choose among and extend, it was necessary early in our research effort to make firm decisions about the character we sought for our NCL models. Our decisions reflected our charter (to improve strategy assessment) and our views about the nature of human decisionmaking. The most important were as follows:

- The decisionmaking of NCL models should be *strategic in character*, while that of the military command levels (recall Fig. 1.2) should be more nearly organizational (or "cybernetic" in the sense Steinbruner uses that term). Figure 2.1 shows some contrasts. However, the wisdom exhibited by alternative NCL models should vary significantly, and should depend upon their perceptions, values, and objectives. Moreover, NCL models should be able to exhibit distinctly flawed behavior of the types we see in the real world of national decisionmaking, however intendedly strategic and wise.
- The NCL models should reflect the decision's logic *process* to some extent, as well as produce reasonable decisions. That is, the logic used by the NCL models should be understandable and comfortable to humans.
- On the other hand, the NCL models should not describe explicitly the various potential interactions among factions. That is, the NCL models should be formally those of *unilateral (monolithic) actors* even though their external behavior might reflect the existence of internal dissension.[1]

Definitions in a Field of Conflicting Terminology

Rationality. Mainstream economists tend to equate rational-analytic thought with optimization of profit or utility (Arrow, 1984:56). We, following Herbert Simon and the more general English-language usage of the term, regard decisionmaking to be rational if merely it is appropriate to the goals sought.

[1]The old joke that a camel is a horse designed by a committee suggests that one can indeed see the impact of internal conflicts without addressing them directly. An NCL model representing a nation with strong bureaucratic conflicts affecting NCL decisions might, for example, exhibit delays, incrementalism, and/or poor integration of actions. Although our techniques could be extended to treat bureaucratic politics explicitly (see, for example, ideas advanced in Kahan, Jones, and Darilek, 1985), we have deferred any such effort because our concern is primarily with external behavior.

Strategic

- Global view
- Reevaluates situation each time (a "state" approach)
- Considers nonincremental options, sometimes new ones
- Designs and/or chooses a plan or process
- Synthesizes information from subordinates, delegates actions, issues orders
- Accepts satisfactory solutions (satisficing) rather than seeking technical optimality

Organizational (Cybernetic)

- Narrow view (e.g., by theater or military function)
- Assesses situation primarily by comparing status with a local predetermined goal and/or position in a process
- Continues a predetermined process or plan, although perhaps taking one or another predetermined branch and perhaps modulating actions within the plan
- Follows a plan, although making some adaptations
- Sends unsynthesized information up the line to superiors (may also synthesize information from subordinates, delegate, and give orders)
- Will frequently attempt to achieve locally optimal performance in particular functions, but will satisfice in others

Characteristics Common to Both

- "Rational" thought
- "Analytic" thought (i.e., use of logic and models)
- Use of heuristics
- Recognition of context and history

Fig. 2.1—Contrasts between strategic and organizational characters

Heuristics. Heuristic solutions are those that are "good enough for the purposes at hand" although not necessarily optimal. Sometimes, however, the term "heuristics" is used to suggest a lower form of reasoning—i.e., "sloppy thinking." In our parlance, heuristic reasoning can be of very high quality and is equally essential to the most brilliant of top-level decisionmakers or to the job foreman making routine marginal adjustments to some industrial process. Because of complexity and uncertainty, real-world decisionmakers *must* use heuristics—this is the problem of bounded rationality discussed by Simon. Thus, heuristic decisionmaking may be highly rational (it may also be emotional and irrational if the heuristics are inappropriate to the goals sought).

Strategic thinking. Our meaning of "strategic" is suggested by Fig. 2.1. Some authors would use "analytic" in this context; others would use "analytic" to imply a dependence on quantitative techniques that we do not want to assume. Note also that our "strategic" decisionmaker may reevaluate the situation each time and thereby be less "plan-bound" than stereotypical organizational models, but he can also be sensitive to history, context, and the cost of changing plans.[2]

Ideal and Nonideal Decisionmakers

With this background of decisions, it proved useful to construct an idealized decisionmaker model as shown in Fig. 2.2 and to itemize as well the types of nonideal behavior we must also be able to reflect (Fig. 2.3).

In developing Fig. 2.2 we drew heavily on work by Simon (1980), Janis and Mann (1977), various books by Peter Drucker (e.g., Drucker, 1974), and our own experience with problems involving command-and-control issues (e.g., Davis, Stan, and Bennett, 1983). This image of an ideal decisionmaker represents an extrapolation of what we observe in good real-world decisionmakers. In that sense, it is a model of human behavior. On the one hand, the characterization focuses on process[3] without specifying the degree to which quantitative analytic techniques, heuristics, or other methods should determine decisions. It does emphasize, however, the importance in the ideal of having a broad view, looking for specially relevant information, and following up on initial decisions (both to assure implementation and to allow for feedback and adaptation in the light of new information).

[2]The NCL model makes decisions based on the current values of a world situation data set, but some variables in that data base can represent important elements of history such as how the conflict began (for example, with Red invading a third country or with both Red and Blue intervening in some third country's civil war). Only in this limited sense, then, can the NCL model have "memory." At a formal level it is Markovian.

[3]The word "process" is another source of common confusion in this field. Organizational (cybernetic) models reflect behavior focused on accomplishing rather cut-and-dried sequences of *actions*—i.e., they are "process oriented." Our strategic models are not process oriented in that sense, but they do represent a plausible and humanlike *decision* process.

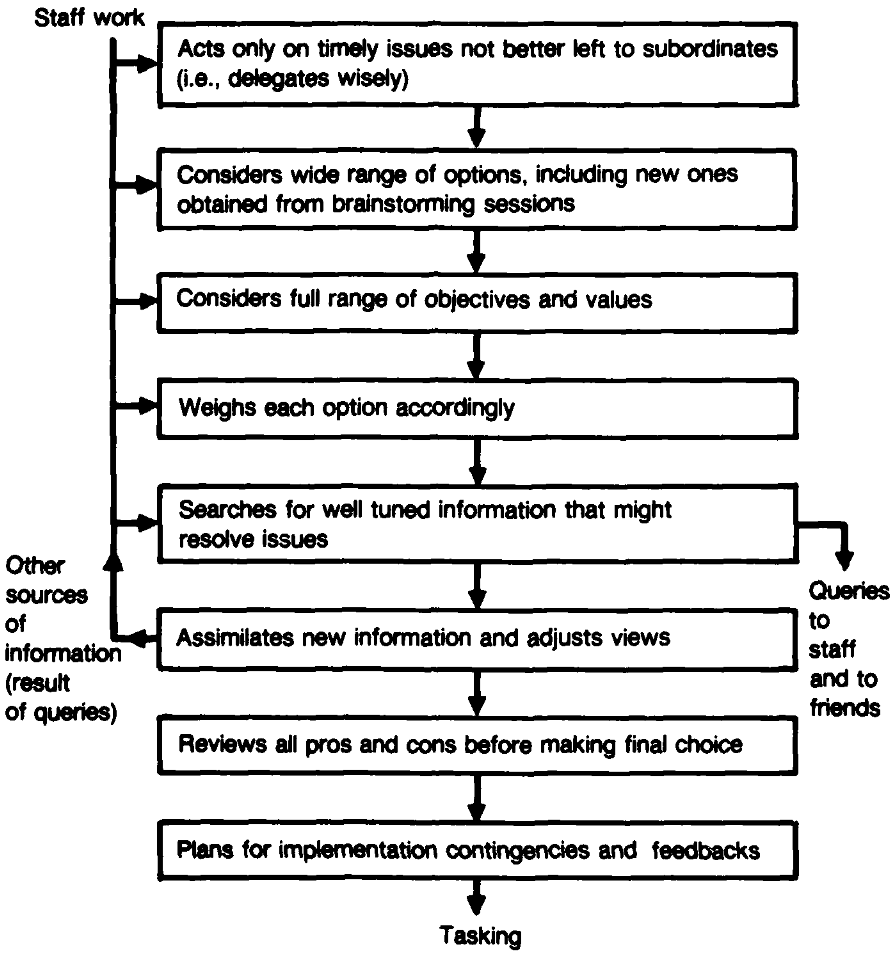

Fig. 2.2—Actions of a mythical ideal decisionmaker

One interesting feature of the idealized decisionmaker is that the image conveyed by Fig. 2.2 is *not* the image of goal-directed search in the sense that phrase is used in much of the technical literature of artificial intelligence and simulation. For example, Fig. 2.2 does not convey an image of the decisionmaker starting out by saying: "This is what I want to do (my goal) and this is how we measure results. Now go out and find me the optimal approach!" To the contrary, the image conveyed by the idealized decisionmaker is one of many conflicting goals, held implicitly in some cases, which must be constantly reexamined and balanced against each other.

As noted above, we must be able to build NCL models that are both wise and that exhibit the typical frailties of decisionmakers. Figure 2.3 summarizes some such frailties, again drawing on the literature of cognitive psychology and other sources (e.g., Janis and Mann, 1977; George, 1980).

The Basic RSAC Process Model for NCL Decisionmaking

Having provided this background describing ideal decisionmaking and common flaws in the behavior of real decisionmakers, let us now discuss the model we have actually adopted. Although this will *not* be immediately evident, our model has much in common with that suggested by Fig. 2.2, as we shall discuss below. Figure 2.4 describes the basic framework of the NCL process model for decisionmaking. The process has three main elements: *situation assessment*, tentative *plan selection* by heuristics, and *plan testing*. It is worth discussing each in turn, although we shall provide more details in Secs. IV and V.

Situation assessment, which is consistent with the strategic character of our model, consists of several subprocesses. First, the model examines the current situation and surrogates for a memory of history; it then uses the opportunity of experience so far in the crisis or conflict to adjust its assumptions about the opponent, third countries, and the laws of war (e.g., best-estimate rates of advance for ground armies). It then uses revised assumptions to project what will probably happen if the current course of action is continued.[4]

The notion of projections is extremely important. In the real world, decisionmakers have staffs. In our work we have a *look-ahead* consisting of a game within a game in which an NCL model projects the future using its own models of reality, which can well be wrong. These look-ahead projections are a mechanism for detecting problems that would not be uncovered by first-order heuristic rules—although, in

[4]As discussed in Secs. III and IV, perceptions of the opponent (Red's model of Blue and vice versa) can be either simplistic or sophisticated. In some instances such as those of a fast moving crisis, we might expect perceptions to be relatively simplistic.

- Fail to delegate adequately or act quickly enough, thereby remaining "behind the power curve."
- *React* constantly to perceived threats or competition rather than insist on seeking the initiative (allow the opponent to make first use of chemical and nuclear weapons).
- Attempt incremental adaptation, even when the probable need for more drastic changes is recognized (continue with a failing strategy or even a failing war, constantly upping the ante).
- See the situation in terms of past personal experience or a particular historical event, even when the analogy is strained (exaggerate the likelihood of the opponent's launching a first strike on the homeland because of sensitivity to the effects of surprise in past wars).
- Insist on masses of information of dubious quality and relevance; suffer effects of saturation.
- Overestimate (or underestimate) the significance of particular risks (preemptively surrender because of exaggerated fear that opponent will use nuclear weapons; or, engage in conventional aggression because of exaggerated confidence that opponent will not use nuclear weapons).
- Reject or ignore complex analyses that cannot be reduced to a few key arguments.
- Ignore important subtleties of complex analyses.

Fig. 2.3—Typical generic shortcomings in actual decisionmaking (examples in parentheses)

retrospect after a look-ahead simulation reveals problems, one can usually construct better heuristic rules that depend only on the current state.

The selection of plans proceeds through a sequence of intermediate decisions once the situation assessment is complete. In the simulation, the available options (analytic war plans) are fixed and the NCL

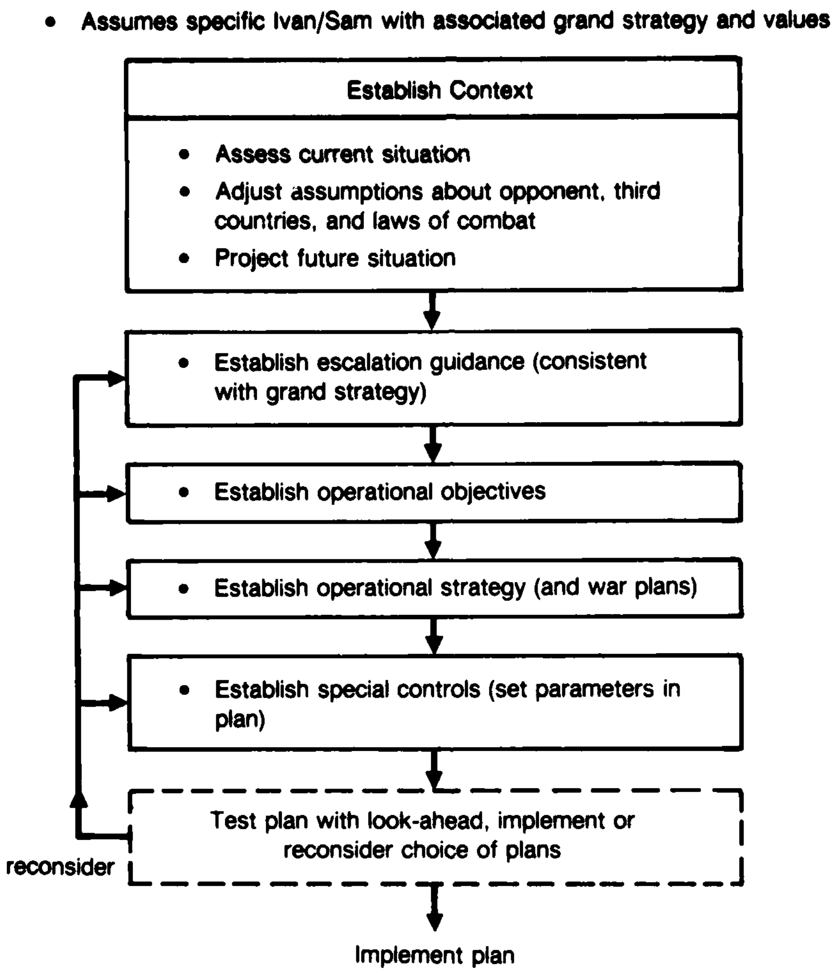

Fig. 2.4—A process model of NCL decisionmaking

model's problem is to pick one. This has the great advantage of allowing us to make complex strategies available to the Red and Blue Agents as inputs rather than attempting to have the agents concoct sensible strategies from whole cloth—something that is far beyond the current state of the art in artificial intelligence except in simple problems (Davis and Winnefeld, 1983). Note, however, that if the NCL models are to be "strategic," then we must include plans that vary substantially from the canonical plans used in most best-estimate work. Controversial plans should be included and the NCL models should be capable of choosing such plans under some circumstances.[5]

As Fig. 2.4 indicates, the choice of plans is limited first by decisions about escalation, then about operational objectives, and finally about operational strategy. Note that the operational objectives differ from grand strategy objectives. A given NCL model represents a particular Ivan or Sam with a particular grand strategy. However, the grand strategy is *implicit* as one looks at the individual rules within the NCL model. Just as it is unlikely that a U.S. President would begin a meeting of the National Security Council by reviewing issues such as the fundamental importance of assuring the survival and continued sovereignty of the United States and its major allies, so also do we not highlight such matters in our decision models. Instead, we focus on making decisions about the scale and extent of conflict and *operational* objectives and strategies. On the other hand, those fundamental decisions are approached very differently by different Ivans and Sams.[6]

A potential criticism of Fig. 2.4 is its sequential character. It is often claimed that human decisionmakers effectively think in parallel about many considerations—looking more at patterns than at logic trains. Moreover, it is evident upon reflection that decisions about escalation, operational objectives, and operational strategy must be in some sense "made together." But this, in fact, is what happens when a human analyst develops rules for each of the decision modules: He may organize the decision process sequentially, but he need not forget his rationale for an earlier decision when visiting a later one. To the contrary, because all of these decisions are supposed to be coherently related to grand strategy, it is in fact *essential* that the various decisions be interrelated and cohesive. In the abstract, this may sound difficult and even esoteric—especially to those attempting to develop modularized generic software independent of substance. In practice, however, it is straightforward.

Once the NCL model has chosen an operational strategy, it has essentially chosen an analytic war plan. What remains is for it to tune

[5]In the real world, it is often difficult for decisionmakers to insist on being presented a good range of options. They must often look outside normal channels to stimulate thinking on the matter. In our work, we must include the unconventional options directly.

[6]We plan to experiment with goal-directed search in the future, but are not sanguine about rapid process *for complex problems*. Indeed, we believe that a major reason for our success in building prototype NCL models has been the decision not to get bogged down attempting to define meta principles for resolving conflicts of principle, but instead to move directly to decision rules for *particular* situations. Even the most thoughtful human policymaker can make individual decisions more easily than he can establish axiomatic principles, and we can hardly expect NCL models to do better at this stage of research.

the plan by setting controls—e.g., controls on rules of engagement, delegation, and use of nuclear weapons.

At this point, then, the NCL model has picked and refined a tentative plan. However, as mentioned above, we conduct look-ahead tests to see whether a plan chosen with heuristics will hold up under greater scrutiny. If it does not, the NCL will consider another plan instead. Note that the failure of a plan in a look-ahead test provides *feedback*, as the arrows in Fig. 2.4 suggest. The NCL model uses feedback information and examines alternative strategies, objectives, and escalation guidance (in that order) until it finds a plan that succeeds on look-ahead. How many different strategies are tested before considering different objectives and how many and which objectives are tested before considering different escalation levels are functions of the Ivan or Sam.[7]

Relationships with Idealized and Less Competent Decisionmakers

Let us now relate the model of the paper to that suggested by Fig. 2.2. At first glance, they are quite different except that both focus on *process*. The differences, however, relate to the technical issues of building an operational model and computer program rather than to the underlying concepts. We make explicit some actions that are implicit in Fig. 2.2 (e.g., situation assessment). And, for the reasons mentioned above, we develop the NCL model's options (analytic war plans) separately so that in the running of our model they are inputs rather than outputs. Nonetheless, **although our representations of decisionmaking are different, the NCL models incorporate the idealized decisionmaker's principal features, such as consideration of multiple options, search for new information, reassessment, and use of heuristic rules in coping with conflicting goals**. At the same time, a particular Ivan or Sam could represent an altogether incompetent decisionmaker if the detailed decision rules were suitably chosen. As examples of how nonideality can be represented, note that the heuristic rules can be simple or sophisticated, that few or many plans can be taken seriously by the model, that

[7]One can circumvent the apparent order sequence of decisions so that, for example, a particular Ivan may seem externally to set operational objectives first and to decide on escalation level later. This Ivan would set stringent standards for success of his plans and would not relax those objectives until he had considered escalatory options. Thus, the feedback arrows in Fig. 2.4 do not imply that every Ivan and every Sam must consider all possible strategies and objectives before varying escalation guidance.

the plans themselves can vary greatly, and that the tests used on tentatively chosen plans may be simple or complex.[8]

HIERARCHICALLY STRUCTURED SITUATION ASSESSMENT

This subsection begins by discussing the informational needs of decisionmakers and decisionmaking models. It then describes a particular framework for providing that information through a hierarchy of variables.

Analysis Versus Synthesis

If there is a single characteristic that distinguishes good decisionmakers from others who may aspire to but never achieve top-level status, it is the capability to perform *synthesis* as well as analysis. Decisionmakers must not only be able to break problems into smaller discrete pieces for more careful study (one definition of "analysis"), but must also be skilled in moving in the other direction—from the chaos of detail to the orderliness of well-structured integration and ultimate implementation. In doing so, they must necessarily develop abstractions that simplify and aggregate, but capture the essence of lower-level considerations.

The process of integration or synthesis is not identical with top-down decisionmaking as that term might be interpreted literally. If the decisionmaker *imposes* his concepts of abstraction and aggregation on the reality of detail, the forest he sees may turn out to be an illusion. Ideally, the synthesis of information should be from the bottom up, with intermediate integration accomplished by those who truly understand the phenomena they are trying to summarize for the purposes of higher-level discussions. An ideal decisionmaker, then, will have a trusted hierarchical staff of experts who will perform these tasks of integration to provide valid bottom-up abstractions. Decisionmakers approaching this ideal are famous for spot-checking the information they receive by asking more and more detailed questions until they can

[8]The model will permit multiple look-aheads using different assumptions about the opponent; it will also permit the NCL to choose plans after comparing their results in look-ahead simulations. In most instances, we do *not* yet include such sophisticated decision processes—for reasons of both technical complexity and verisimilitude. It should also be emphasized that heuristic rules based solely on current information can be quite good.

judge its basis and also assess the degree to which implicit values used in the assessments are consistent with their own.[9]

All of this suggests an approach to decision modeling that emphasizes a complex formal bottom-up synthesis of information—even if the decisionmaker himself is to be represented as taking a top-down approach in the sense of making high-level strategic decisions that are then passed down into the lower levels of the organization for detailed implementation.

Quantitative and Qualitative Reasoning

Another serious problem for real-world decisionmakers as well as analysts attempting to develop decisionmaking models is the sheer volume of potentially usable information. In traditional quantitative work, analysts have usually reduced the problem to focusing on one or a few measures of effectiveness such as the ratio of hard-target nuclear weapons, the locations of Forward Line of Own Troops (FLOTs), the rates of advance in various key sectors, or the expected damage levels to particular categories of fixed targets. Figure 2.5 provides a somewhat longer list.

• Rate of advance	• Status of naval operations
• Ground position	• Troop morale
• Deliverable nuclear warheads[a]	• Ground force attrition level
• Deliverable hard-target weapons	• Ground force attrition rate
• Status of support forces	• Air force attrition level
• Status of sustainability	• Air force attrition rate
• Status of control structure	• Naval force attrition level
• Status of civil defense	• Naval force attrition rate
• Alliance cohesiveness	• Attrition rate of nuclear weapons

[a]Equivalent megatonnage, EMT.

Fig. 2.5—A standard list of variables used to evaluate a political-military situation

[9]Realistically, of course, there can be so much cognitive dissonance between those at the top and bottom that the decisionmaker or his personal staff must perform the simplifications without adequate knowledge of lower-level events.

One problem with such lists is that the variables measure static capabilities rather than functional capabilities—although, to be sure, the two are related. Figure 2.6 lists measures that are more functionally oriented and correspondingly more directly relevant to sound decisionmaking. By contrast with common assumptions in strategic nuclear studies, decisionmaker judgments will often depend far more on assessments of absolute functional capability than on the ratios of nuclear weapons or other measures dear to the hearts of quantitative analysts.

Figure 2.7 provides yet another list, this one more qualitative yet, but closely related to the topics we are trying to treat. Indeed, as argued in Davis and Stan (1984), humans acting in political-military war games tend to focus on these issues when making the cosmic decisions about escalation and termination. In many cases, what is at issue are *symbols*, symbols of restraint or of escalation, and symbols of the crossing of tacit boundaries (about which the two sides may or may not have a common sense). These notions owe much to Thomas Schelling, Herman Kahn, William Kaufman, and others who plowed

- Prospects for completing an ongoing offensive
- Prospects for defeating the enemy's campaign
- Capability in a nuclear strike to reduce enemy nuclear offensive capabilities to low levels
- Capability to intercept a large fraction of an attacker's force and to limit damage significantly
- Capability to ride out a nuclear attack and still retaliate effectively
- Capability to control events in the homeland and to conduct effective civil defense and continued military operations
- Capability of naval forces to deny enemy use of particular regions

Fig. 2.6—Some representative measures of functional capability

- History (e.g., origins of conflict)
- The apparent nature of the opponent
- Strategic or tactical warning of opponent escalation
- Success or projected success in meeting objectives
- Costs and projected costs of alternative actions
- The sanctity of the homelands
- The sanctity of strategic forces (and space?)
- Survivability of command-control and critical forces
- The opponent's use, even if for nonstrategic purposes, of systems considered "strategic" (in the nuclear sense)
- Size and nature of naval warfare

Fig. 2.7—A list of key symbolic iss· ·

the war gaming fields back in the 1950s. Also important, however, are such issues as "who" one thinks the enemy is and how one got into the current situation.

Continuing with this theme, real-world decisionmakers tend to have very different perspectives than traditional quantitatively oriented strategic analysts. In particular, they concern themselves with a broad range of *qualitative* characterizations. In ordinary crises, these would include prominently judgments about political factors, the detailed personalities of other world figures and their advisors, and recent political intelligence. For example, decisions in ordinary crises often reflect what leaders feel they must do to maintain power and legitimacy. In extraordinary crisis and conflict, however, we hypothesize for the purposes of analysis that the decisionmakers would focus less on personality factors and more on high-level political-military considerations, which would include an assessment of likely opponent behavior but with less attention paid to the mechanisms generating that behavior.

A similar phenomenon occurs in strategic planning in any organization. Although *operating* decisions may take into account all sorts of detailed low-level information valid at the time of decision, strategic decisions should take a longer view and depend more on enduring realities (or on secular trends) than on transient factors such as the current managerial chaos of a competitor resulting in temporary advantages.

Strategic decisions should also transcend the petty quarrels and competition within the organization.

Policy-Level Assessments and Value Tradeoffs

A fundamental characteristic of high-level decisionmaking is that it is replete with value-laden tradeoffs. So, for example, in assessing the current situation the decisionmaker must somehow fold together the various pieces of information and come up with an overall judgment: "Well, basically, our status in the European war is pretty good—we're holding our own and time is working with us at this point." In reaching this assessment, he has necessarily made a tradeoff between "apples and oranges."[10]

To illustrate this, suppose that the situation on the ground in Europe is mixed from the Soviet point of view: Soviet forces have made progress on all fronts, but more slowly and at higher casualty rates than expected. In reaching an overall assessment of the ground war's status, how should one balance off these several factors? A traditional quantitative analyst might argue that it is really a matter of finding the right weighted sum of measures of effectiveness. Although that approach may be comfortable for traditional quantitative analysts, it can do violence to the problem of situation assessment and decisionmaking, in part because it seldom seems natural to decisionmakers to work with weighted sums of different measures: They would rather do the tradeoffs more intuitively using details of context that would be ignored in a typical weighted-sum approach (e.g., reports of an imminent enemy collapse or loss of a key general).[11]

Pursuing the example, it is clear that there is an important tradeoff between current position and the attrition suffered in arriving at the current position. Although attrition might be a secondary factor in most instances, sufficiently great attrition would call into question prospects for continued gains and might violate human values or leave the commander vulnerable to criticism and even replacement.

To cite another example, consider the tradeoff between status of the ground war with respect to such items as troop locations, attrition,

[10]For a discussion of tradeoff issues within the research domain of multiattribute decision analysis, see Keeney and Raiffa (1976), which tends to emphasize quantitative approaches involving utility theory. See also Kahan (1979) for a discussion of alternative tradeoff procedures.

[11]Soviet military practice makes much greater use than does American practice of algorithmic techniques for reaching operational-level decisions and tactical-level decisions. The Soviet attitude seems to be that hastening those decisions is extremely important and that maintaining troop control requires regularizing the decision process (see, for example, Hemsley, 1982).

rates of movement, and status of nuclear forces in the same theater. Although the ground war's status might seem to be the dominating factor in an overall assessment, that would not obviously be the case if the balance of nuclear forces had changed dramatically, leaving one side with the capability to escalate into a domain in which the other side had minimal capabilities (escalation dominance). Again, a tradeoff is necessary to characterize the overall status of the war in that theater.

In subsequent sections we shall provide concrete rule-based examples of such tradeoffs. Our point here is that making such tradeoffs is an essential part of building NCL models and that the tradeoffs inherently involve values and perceptions as well as purely objective factors. Thus, alternative Red and Blue Agents will differ not only in their decisions about escalation, objectives, and strategy in a given situation described in high-level terms, but also in how they characterize the same situation.

The Hierarchy of Variables

With this background of observations, let us now describe the particular hierarchy of variables we have used in our prototype NCL models. This hierarchy constitutes a reasonably generic *framework* within which diverse particular NCL models can readily be developed.

The image we have attempted to use in our prototype models is one of the NCL focusing on three fundamental aspects of the current situation: Status, Prospects, and Risks:

- *Status*: "How are things going in the main theater? What about elsewhere?"
- *Prospects*: "What's the best-estimate projection if we continue on the current course of action (current analytic war plan)? What about for the key options?"
- *Risks*: "How likely is it that events may unfold very differently than under best-estimate assumptions? How likely is it that the opponent will escalate? If such things happen, how serious would the consequences be?"

The answers to such questions are naturally qualitative and subjective as in, "Well, sir, the risks are high and our prospects are only marginal." If such characterizations are indeed the natural currency of high-level discussions, as we believe they are, then it is appropriate to use those high-level concepts directly in our decision rules. The key, of course, is that we must also specify the way in which it is determined that risks are "high." Deferring that for a moment, let us complete our image of the decision process. The decisionmaker, having been deluged with information in the form of briefings, memoranda, personal advisories, and so on, reaches his decision finally with something like:

> Well, gentlemen, if I understand what you have been telling me these last few hours, and if I try to patch together some of the pieces that came out one at a time in our meeting, then it seems that our current situation is pretty good—we have achieved our principal objectives, although not everything we had hoped for. We could push on, but prospects for further progress appear only marginal and there are big risks. If that is all correct, then I conclude we should begin to consolidate our gains and wind down our actions. Do you agree?

Let us emphasize at this point that high-level concepts such as *Situation* (i.e., aggregated concepts appropriate to a top-down view) can be extraordinarily complex. For example, they can involve not only objective factors such as military gains or profits earned, but also subjective factors such as political perceptions of victory.

To develop this further, then, we must relate the high-level concepts explicitly to lower-level concepts that are increasingly more specific and objective.

Figure 2.8 illustrates the hierarchical manner in which we can relate the high-level concepts to the more detailed events in a war game simulation. For example, the high-level concept of *Status* is a function of the intermediate-level variables *Main-theater Status*, *Other Status*, *Opportunity*, and *Warning*. These, in turn, have values dependent upon more detailed variables such as the status of the ground war (which is measured in terms of FLOT positions and attrition) and force status. The higher-level variables in Fig. 2.8 are intended to be generally applicable, but the lowest-level variables are simplifications adopted for the sake of prototyping. For example, there is no mention of issues such as the status of civil defense or homeland mobilization.

Some definitions may be useful in viewing Fig. 2.8:

- Main-theater status summarizes overall progress in the main theaters of action.
- Other-status is a composite assessment of how things stand "elsewhere" than in the main theaters (e.g., are they worse than expected?).
- Opportunity is a measure of whether one should raise objectives (is there a vacuum somewhere, or have the correlations of force changed dramatically?).
- Warning means what it suggests.

Status	Main-theater status	Ground status	FLOTs
			Attrition
		Force status	Force levels, ratios
			Attrition
		Political status	Alliance
			Other
	Opportunity		
	Other status	Level of conflict	
		Other	
	Warning (urgent and nonurgent)	Strategic warning	
		Tactical warning	
		Projected warning	
		Critical intelligence warning	
		Ultimatum warning	

Prospects	Main-theater status	From force levels, ratios
		From simple projections
		From look-aheads
	Other status	From force levels, ratios
		From simple projections
		From look-aheads
	Opportunity	

Risks	Escalation	From force levels, ratios
		From simple projections
		From look-aheads
	Main-theater status	From force levels, ratios
		From simple projections
		From look-aheads
	Other status	From force levels, ratios
		From simple projections
		From look-aheads

Fig. 2.8—Simplified view of hierarchical determinants of situation and prospects (the top line means that the FLOTs help determine ground status, which helps determine main-theater status, which helps determine status)

The variables on the far right of this figure are the ones that can be observed objectively in the simulation or in the real world; those to the left are increasingly abstractions that depend upon values, judgments, and perceptions.

To recapitulate, the logic behind these hierarchically related variables is essentially as follows:

- A decisionmaker being provided a situation assessment needs to know something about how things are *now* and how things are likely to be *in the future*. Furthermore, he needs to have not only a best estimate of such matters, but also some estimates that are more conservative (and, perhaps, more optimistic). That is, he needs to know not only how his staff thinks things will *probably* go (prospects), but also how they *might* go in a bad "roll of the dice" (risks).
- In making predictions, it is useful to distinguish between predictions based on: (1) current status and straightforward heuristics (e.g, given a 4:1 force ratio, a commander might be optimistic without doing any detailed calculations); (2) formal projections on disconnected parts of the overall problem (e.g., projecting the arrival of enemy reserves assuming no basic changes of enemy strategy such as shifting forces from other theaters); and (3) formal look-ahead projections simulating everything one can, including national political decisions and reactions to them.

There are other principles involved, but these are enough to motivate the general structure of Fig. 2.8. Obviously, there are alternative constructions and in our applications work we are adding significant complexity. However, our initial work suggests that the basic approach is rather flexible, as intended: The enhancements necessary are primarily those to lower levels of the hierarchy of variables. Thus, although it is assuredly feasible for others to use different names or concepts in developing NCL models, it is also possible to use ours to reflect quite a range of plausible behavior patterns. This is important practically, because building operational computer programs is a difficult and time-consuming process. To the extent that one can hold the structure constant and address different problems by changing only the "data" (in our case, the detailed rules establishing the values of the several layers of variables, and for making the decisions indicated in Fig. 2.8 in terms of the higher-level concepts such as prospects, risks, and warning), one can work with much greater efficiency and focus on substance rather than semantic finepoints.

It was our intention that the hierarchies in Fig. 2.8 be sufficient to accommodate a broad range of alternative Ivans and Sams. That is, the intention is that different NCL models will have in common the key variables identified above, but will differ in the way they evaluate the variables. So, for example, all Ivans will recognize that risk is an issue, but some may weight it relatively low when it comes time to characterize the current and projected situation.

III. CHARACTERIZING IVANS AND SAMS

INTRODUCTION

Our general approach to NCL modeling depends on the notion of *alternative* models of the United States and Soviet Union (Ivans and Sams). Much of this report describes the framework for a generic computer program that will accommodate many such alternative Ivans and Sams. In practical terms that means that if an analyst wishes to develop a different model of the Soviet Union, then—within broad limits—we can probably accommodate it if he will translate his mental model into decision rules for the computer program. However, for that to be feasible the analyst must have some conception about how to write those rules—i.e., how to sharpen his image of the Soviet Union so that it will indeed translate into coherent and sensible rules.

This section, then, describes a methodology for developing and sharpening an image of the United States or Soviet Union *before* attempting to write the first decision rule. The methodology is straightforward and has been used by all of the authors separately to develop a range of prototype Ivans and Sams. Without going through something like the exercise we suggest here, it would be difficult to produce decision rules that were either coherent or in any sense comprehensive.

The methodology consists of six discrete steps, each of which we will illustrate with one Ivan and one Sam. The steps are all useful, but their relative usefulness is subjective.

1. *Limit the problem area* to some class of scenarios, however large.[1]
2. Develop a *prose description* of the Ivan or Sam—i.e., an essay.
3. Fill out a *temperament checklist* that characterizes Ivan or Sam in terms of formal attributes of the agent's political, attitudinal, and perceptual personality.
4. Construct a series of notional *grand strategy decision trees* representing graphically how the Ivan or Sam might look at the current strategic situation and the decision points he is likely to have in the course of the campaign.

[1]It is impractical to develop altogether comprehensive decision rules, especially at this stage of experience. Furthermore, it is important for the analyst to understand explicitly for what cases his model will not be applicable. Indeed, that information can be made part of the computer program and can be used to warn unwary users.

5. Fill out a *transition matrix* for the most likely escalations and deescalations that would be entertained by the Ivan or Sam.
6. Establish some *guiding principles.*

We shall define each of these by example in what follows.

LIMITING THE PROBLEM DOMAIN

Most strategic nuclear analyses treat very few scenarios and treat those in the simplest of terms—e.g., day-to-day alert and generated alert as the "scenarios" for a single massive exchange of nuclear weapons. Consistent with general goals of RSAC research (Davis and Winnefeld, 1983), we consider much richer classes of scenarios extending from large-scale crises through general nuclear war. Nonetheless, in developing our prototype models we have limited the class of scenarios enough to give us a foothold. By and large, most of our work has assumed scenarios in which (Davis and Stan, 1984):

- War begins with the Soviet invasion of Southwest Asia. No nuclear weapons are used there until and unless they are used in Europe. There may or may not be Soviet threats in Europe (e.g., a Berlin blockade or mobilization).
- War may (or may not) spread to the naval theaters worldwide and to Europe.
- War in Europe may or may not escalate to nuclear war or to intercontinental nuclear war.
- Nothing dramatic happens in "other theaters" such as China and Africa, although such theaters represent a drain on forces and establish constraints and concerns within which the superpowers must act.

The Soviets write about analogous scenario classes except that war begins with aggression by the United States.

There are enormous variations possible within this general class (e.g., variations in warning, alliance cohesiveness, the time gap between activities in the several theaters, and Soviet strategies), but the limitations are very useful. For example, they suggest a particular way to focus on a moderate number of theaters and conflict levels. These in turn provide a strategic structure for situation assessment and other functions. We shall not discuss that in detail in this report (see, however, Davis and Stan, 1984). Figure 3.1 shows an aggregated escalation ladder of situations useful for first-order descriptions across the scenario class discussed above. Figure 3.2 shows how one can display a

1.	Prep	Soviet preparation for invasion of Iran
2.	Term	Termination of war
3.	SWA-Sov-conv	Soviet invasion of Iran
4.	SWA-Sov-conv/Eur	Invasion of Iran plus a Berlin blockade or other demonstrative action in Europe
5.	SWA-gen-conv	U.S.-Soviet war in SWA; fought with conventional weapons only
6.	Eur-demo-conv	Demonstrative conventional actions in Europe (e.g., Soviet blockade of Berlin, possibly with some combat)
7.	Eur-gen-conv	General conventional war in Europe
8.	Pnuc-Eur-gen-conv	General conventional war in Europe after use of nuclear weapons there (assumes 24 hours of no nuclear use)
9.	PIC-Eur-gen-conv	General conventional war in Europe after intercontinental nuclear exchange (assumes 24 hours of no nuclear use)
10.	Eur-demo-tac-nuc	Demonstrative operational-tactical nuclear war in Europe
11.	Eur-gen-tac-nuc	General operational-tactical nuclear war in Europe
12.	PIC-Eur-gen-tac-nuc	General operational-tactical nuclear war in Europe after an intercontinental phase (assumes 24 hours of no intercontinental strikes)
13.	Eur-demo-strat	Demonstrative strategic nuclear war in Europe
14.	Eur-CF-strat	Counterforce strategic nuclear war in Europe
15.	Eur-strat	Strategic nuclear war in Europe
16.	IC-CF-strat	Intercontinental strategic counterforce war
17.	IC-strat	Intercontinental strategic nuclear war

Fig. 3.1—A simplified ladder of multitheater conflict situations

Conflict Level	SWA Chem? Yes	SWA Chem? No	Other Naval	Europe Chem? Yes	Europe Chem? No	Europe USSR Home? Yes	Europe USSR Home? No	Space Tac? Yes	Space Tac? No	Other Land	Intercont. (U.S./SU)
General Strategic											
Counterforce Strategic											
Demonstrative Strategic											
Post IC General Tac. Nuc.											
General Tactical Nuclear	RB		RB	RB							
Demonstrative Tactical Nuc.						B					
Post Intercont. Conven.											
Post Tac.-Nuc. Conven.											
General Conventional	RB		RB	RB				RB			
Major Nonnuc. UCW											
Demonstrative Conv.											
One-Superpower Conv.											

NOTES:

1. Red and Blue are engaging in conventional, chemical, and tactical nuclear warfare in Europe and Southwest Asia. Blue has launched a demonstrative nuclear strike into the Soviet Union, in the context of the European war. Both sides are attacking satellites, but not strategic satellites related to missile warning.

2. UCW = unconventional warfare or planned disturbances in cities and other rear areas (e.g., organized strikes).

Fig. 3.2—A simplified representative conflict state

top-level view of worldwide conflict with somewhat more discrimination. Later, in Secs. IV and V, we shall discuss some of the more important lower-level variables as well.[2]

The point here is that the analyst wishing to develop an Ivan or Sam rule-set should consciously worry about the class of scenarios for which his rules are to apply, and should worry as well about which theaters and levels of conflict would constitute a useful representation. As merely one illustration of how important this is, note that Figs. 3.1 and 3.2 do *not* deal well with classes of wars starting in the Middle East with the United States and Soviet Union both sending forces to support factions of a regional dispute. That would require a somewhat different representation. Rule-writing would also be more complicated because, for example, it would be quite plausible for both superpowers to believe that they were reacting to the other's aggression; moreover, both superpowers might feel trapped by unexpected events. By contrast, the class of scenarios implicit in Fig. 3.1 assumes initial Soviet aggression (however motivated and rationalized).

PROSE DESCRIPTIONS

The next step is to construct a short informal description of Ivan or Sam in English prose. It should be both action-oriented, in the sense of describing *how* the Ivan or Sam would behave, and motive-oriented in encapsulating *why* the agent would engage in such behavior. The purpose of describing "how" is to portray the agent in concrete terms subject to critique. The purpose of describing "why" is to provide a guideline so that the rule-writer and readers may generalize beyond the concrete instances to new situations. Let us now consider examples for Ivan K and Sam 5. *Neither verbal description is intended to represent a best estimate of actual U.S. or Soviet behavior*; rather, they are merely illustrative.

---

[2]The concept of escalation ladders is old, but in practice it has serious flaws because "level of conflict" is multidimensional. Figure 3.2 and more complex analogs are preferable for communications and careful thought (Davis and Stan, 1984). Ladders such as that in Fig. 3.1 are still useful analytically (e.g., for grouping classes of situation) but must be accompanied by many detailed definitions, since the order of rungs becomes somewhat arbitrary. For example, it is awkward with Fig. 3.1 to treat states such as (1) war in SWA accompanied by large-scale horizontal escalation; (2) conventional war in Europe coupled with nuclear war at sea; or (3) conventional war in Europe coupled with extensive antisatellite activity. In future work we will be using more complex structures that ought not be considered "ladders" in the usual sense of each rung having an objective order.

Verbal Description of Ivan K

Ivan K is neither simplemindedly doctrinal nor softheadedly optimistic. He tends to be somewhat aggressive, risk-taking, and contemptuous of the United States, but is ambivalent about the latter, believing that the United States and NATO, if provoked to war, would tend to become aggressive and irrational. Ivan K is "very Russian" in his belief that basic Soviet military doctrine is essentially correct, although not always applicable. He is *strongly* averse to nuclear war. While he would be loathe to instigate a general nuclear war, he believes that any war with the West might well escalate and that the Soviet Union could survive and, in some sense, "prevail."

If Ivan K sees opportunities for expansion in Southwest Asia or other areas, he will capitalize on those opportunities and make plans for invasions. If expansion plans are successful, he will revise his goals to include greater expansion. For example, if a plan to acquire Iran were successful, Ivan K would expand his goals to acquire more of the Persian Gulf region.[3]

On the other hand, to minimize direct conflict with the United States, Ivan K *might* be willing—temporarily—to moderate his goals of occupying Iran to settle for an endpoint with Soviet occupation of northern Iran (or all of Iran), and the United States out of SWA altogether. As a part of his grand strategy for a SWA campaign, Ivan K would plan actions elsewhere in the world to divert the United States and divide its alliances.

Ivan K has absolutely no intention of starting a real war in Europe, much less allowing that war to become nuclear or general. However, should Ivan K find himself "having to" invade Europe (something remote from his thinking at the war's outset), he would attempt to keep the war short, limited, and decisive. Should Ivan K find himself fighting a nuclear war in Europe, he would attempt to keep it limited (although he might consider escalation to nuclear war at sea in response to some NATO actions there or on the European continent). However, he would concern himself with the dangers of escalation and would be skeptical about the feasibility of fighting a theater-strategic conflict.[4]

---

[3]Ivan K's military leaders might argue from the outset for not involving themselves in the Persian Gulf except to occupy the entire region. Ivan K, however, would be more incrementalist.

[4]One reviewer commented at this point that Ivan K is a "mealy-mouthed, middle-of-the-roadish Ivan which may satisfy the grammar of simulation and the consensus seeking nature of group analysis. But it is also unlike any of the real Ivans who, in my judgment, we would encounter in crisis. . . ." Other reviewers have found Ivan K very realistic. This demonstrates the *necessity* of having alternative models of the Soviet Union and reminds us that NCL models are simply not "expert systems" in the usual sense of that term—there are no experts in this domain, only specialists with varied opinions.

Finally, although Ivan K does not seek general nuclear war, he would not shrink from one once he believed it inevitable. Should he find himself in a war that he believed the West to have started in an attempt to dissolve the Soviet empire, or if he found himself in a war in which the West seemed to have taken on such ambitions, he might well set a grand strategy aimed at the effective destruction of the United States as a world power. In such a case, he might be willing to escalate even to general nuclear war.[5]

Verbal Description of Sam 5

Sam 5 represents one embodiment of "flexible response" U.S. policy, *with respect to NATO and the European theater.* This Sam will escalate, if at all, in accordance with a theory of maintaining or reestablishing deterrence through limited, demonstrative uses of nuclear weapons. Put another way, *Sam 5 views nuclear weapons as war preventing rather than war fighting tools.* Therefore, Sam 5 does not necessarily act for militarily decisive effect, but instead attempts to communicate resolve. Indeed, if demonstrative nuclear attacks escalate into general tactical nuclear warfare, Sam 5 regards his strategy as having been unsuccessful. Although his and his allies' peacetime policies might imply strategic nuclear use if deterrence were not reestablished by initial nuclear use and if the war were going badly, his actual behavior would depend sensitively upon circumstances.

Sam 5 believes in quick response, given support by his allies (which he initially believes to be forthcoming), and will venture forth with what he believes to be a deterrent force. He sees Iran itself as more important for the allies than his own immediate self-interest but important for U.S. self-interest as the first domino on the Gulf. The grand strategy for Sam in part depends on his allies. If they join with him in making a show of strength in Iran, he will defend it strongly, albeit not using nuclear weapons. Any action in Europe, either a blockade of Berlin or moves into the Federal Republic of Germany and Denmark, will be met by prompt and decisive action, including tactical nuclear weapons, if necessary. Intercontinental strategic weapons would be used only as a last resort, and no homeland strikes against the Soviet Union would be planned unless there were tactical warning of intercontinental strategic nuclear war or strategic nuclear attacks in Europe. (Other Sams would launch an intercontinental strike if NATO's defenses were about to collapse and tactical nuclear weapons had failed to stop the Soviet advance.)

---

[5]To illustrate how the essays cannot be complete, note that this description says nothing about Ivan K's attitudes about war at sea. We could add some comments, but other "holes" would remain. How much is enough depends on the applications for which the model is to be used.

TEMPERAMENT CHECKLIST

The next step in characterizing Ivans and Sams in preparation for rule-writing is an attempt to characterize the agent's temperament using a formal checklist.[6] If the essays are systematically written, then filling out the checklist will be easy. Conversely, if one uses the checklist in writing the essays, then the two efforts converge. An agent *temperament* describes in list format a major agent's (Red or Blue) general orientation toward political-military policymaking. The list consists of *attributes* or dimensions along which a general orientation may be categorized. By having to think in terms of formalized attributes, the rule-writer is reminded of other issues, which may also cause him to reconsider and rewrite the verbal description. In future work we expect to use the attributes directly in NCL programs, but for now they are purely a rule-writing and communications aid.

Temperament is conceptualized in two different levels of detail (20 and five attributes, respectively). The first gives a systematic description, the second an abbreviated sketch.

The Four Themes of Temperament

Temperament is defined as four themes that constitute an agent's cognitive structure. Each theme contains a number of attributes. The four themes are: Flexibility, Political Orientation, War fighting Style, and Perception. Each attribute may take on a number of values:

1. Flexibility attributes dictate the extent to which the agent will be able to change his goals and his perceptions of other agents. There are two flexibility attributes:

- *Flexibility of Perception* deals with the ease with which an agent can change his perceptions of the nature of another agent, as information is obtained about that other agent's behavior. The values that flexibility of perception can assume are: rigid (unwilling to change perceptions), conservative (slow to change perceptions), Bayesian (inclined to update perceptions by balancing new information against original conceptions), and ahistorical (inclined to revise perceptions quickly in the light of new information, even to the extent of ignoring previous information).
- *Flexibility of Objectives* deals with the ease with which an agent can change his general objectives. This might involve changing grand strategy in light of circumstances, or altering the general

[6] See Kahan, Schwabe, and Davis (1985) for more details.

objectives that are designed in the service of that grand strategy. Flexibility of objectives can take on the values: flexible, limit-setting, and resolute.

2. Political attributes represent the political orientation of an agent with regard to nuclear policy, depth of commitment to defend various national interests, and attitude regarding the agent's role in world affairs. These attributes describe relationships with other nations, as well as policy in the use of nuclear weapons, and political (as opposed to military) constraints on actions. There are four attributes of political orientation:

- *Use of Nuclear Weapons* describes the actual national policy governing the circumstances, if any, under which nuclear weapons might be employed. This is a guideline policy, which does not necessarily determine actual behavior but rather sets a baseline stance. The values for this attribute are finely distinguished, and comprise: never use, never use tactical nuclear weapons, no first use/respond in kind, no first use/respond to win, preempt, first use to maintain the initiative, first use for defense only, and do not treat nuclear weapons as qualitatively different from conventional weapons.[7]
- *Commitment* describes the extent to which an agent is committed to the defense of other countries or regions, as a function of their importance to national interests (homeland, vital national interest, intrinsic national interest, indirect interest, and diffuse interest). For each of these levels, the possible levels of commitment are: Do not commit forces, use conventional weapons but not at high levels, use whatever conventional force is appropriate, use nuclear force only if conventional weapons fail, and use nuclear force to win.
- *Expansionism* describes a major agent's attitude toward his own empire-building ambitions. This particular attribute is of primary importance for Ivan; Sam is taken as a status-quo power in all of his incarnations, although this need not be so in principle. The values for expansionism are: status-quo, conservative, opportunistic, and adventurous.
- *Unilateralism* describes the extent to which a major agent consults relevant allies before acting. This attribute is primarily

[7]Attitudes about nuclear war at sea may be special and will be considered separately in future work. One interesting special issue is whether one would use conventional weapons aggressively to reduce the opponent's nuclear capabilities. We assume here the answer is yes. Another issue not addressed here is acceptance of sanctuaries.

important in describing various Sams; Ivan is assumed to demand autonomy in all of his incarnations. The values here are: autonomous, consultative, protective, and subservient.

Attached to each political attribute is a *priority rating* that indicates the centrality of that particular attribute in the agent's cognitive decisionmaking structure. The higher the priority rating for an attribute, the more an agent is constrained to choose a course of action consistent with the attribute, even if situational elements and other attributes indicate the desirability of alternative behaviors. If two political attributes prescribe contradictory courses of action, the priority rating determines which will have the greater influence on the decision. For example, a situation might arise in which a particular area is severely threatened by an enemy. The Commitment attribute might prescribe a nuclear defense of this territory, but the Use of Nuclear Weapons attribute might prohibit a first use strike. If the Commitment attribute has a higher priority, then a war plan employing nuclear weapons will be tested, while if the Use of Nuclear Weapons attribute is of higher priority, then conventional alternatives will be pursued.

3. War fighting style has four attributes, which also have priority ratings. The four are:

- *Risk Proclivity*, or the degree to which the NCL is willing to attempt to achieve desirable outcomes that have low probabilities of success or high costs if the attempt fails. The three values for this attribute are: risk-taking, pragmatic, and risk-averse.
- *Operational Daring*, or the degree to which plans that deviate from traditional military doctrine will be entertained. This attribute has the values: daring, open, and standard.
- *Insistence on Initiative*, or the perceived need to dictate the level and extent of conflict as opposed to reacting to opponent behaviors. Initiative here refers to actual military engagement, not to preparation. The values are: proactive, preemptive, and reactive. Proactive means that the NCL will initiate action even in the absence of a tactical threat, if necessary to maintain momentum.
- *Look-Ahead Tendencies*, or the degree to which military planners pursue the potential consequence of proposed actions by using simulation techniques anticipating outcomes and reactions of the opponent. Values for look-ahead tendencies are: shallow, moderate, and deep.

4. **Perceptual attributes** are a major agent's view of the other agent's attributes for political orientation and war fighting style.[8] Just as there are priority ratings attached to the political and war fighting attributes to determine their importance in agent planning, so also are there perceived priority ratings to predict how the other agent will behave. In this sense, the perception attributes mirror in procedure, if not in content, the political and war fighting attributes and take on the same values. There is no necessary correspondence between one agent's perception of the values of the other agent's attributes and that other agent's real values. For example, a Sam might believe that a particular Ivan is willing to take substantial risks when that Ivan is in fact fairly conservative. Or, an Ivan might misperceive Sam as expansionistic instead of oriented toward the status quo.

In addition to perception attributes for the other major agent, there are two attributes for the perception of the anticipated participation of the agent's own allies and the opponent's allies. These third-party attributes follow the behavior attributes developed for Green Agent (Schwabe and Jamison, 1982). There are, therefore, ten perception attributes, as follows:

- *Perception of Opponent Use of Nuclear Weapons*
- *Perception of Opponent Commitment*
- *Perception of Opponent Expansionism*
- *Perception of Opponent Unilateralism*
- *Perception of Opponent Risk Proclivity*
- *Perception of Opponent Operational Daring*
- *Perception of Opponent Insistence on Initiative*
- *Perception of Opponent Look-Ahead Tendencies*
- *Perception of Own Ally Participation,* or the degree to which an agent's allies (largely NATO for Sam and the Warsaw Pact countries for Ivan) are committed to participating in planned actions. The values for this attribute follow Schwabe and Jamison (1982) and are: noncoordinate, coordinate, noncombatant, cobelligerent, and nuclear-releaser.
- *Perception of Opponent Ally Participation,* or the corresponding perception of the opponent's allies' participation. The values for this attribute are the same as for the one immediately above.

[8]It will be important in future work to consider explicitly perceptions of the opponent's perceptual flexibility.

Representative Checklists

Figures 3.3–3.5 present examples of temperament for Ivan K and Sam 5, following their respective verbal descriptions presented above. Figure 3.3 presents the short list of attributes for both agents, Fig. 3.4 outlines the temperament of Ivan K, and Fig. 3.5 gives the temperament of Sam 5. In each case, a full menu of attribute values is shown, with the selected value indicated.

GRAND STRATEGY DECISION TREES

As mentioned repeatedly, a given Ivan or Sam should embody a grand strategy. But what does that mean and how do we make the notion concrete? In reality, doing so is complex if our standards are too high—grand strategies are seldom clear-cut except in retrospect for the winners. Furthermore, it is not uncommon for grand strategies to be discussed without adequately dealing with the branch points—the "What ifs."[9]

To cope with this problem we have found it very useful to have blackboard sessions in which we move from pure abstractions to relatively concrete views of the world as it might be seen by the Ivan or Sam in question. These views take the form of decision trees—not "complete" trees covering all the possibilities, or even all possibilities within the RSAC simulation—but rather, trees indicating where the agent thinks he is going and what some of the more important contingency branches are that he recognizes consciously. To be sure, it would be unusual to find a real-world head of state willing to lay out a grand strategy decision tree on a blackboard for his staff, but we can nonetheless discuss what that tree might be like.

Once again it becomes important to limit the range of scenarios. It is impractical to lay out some type of multidimensional tree covering all the theaters and all the possible combinations of events—that, indeed, is a better job for computer code constructed from decision rules. With this in mind, then, let us now consider some simplified grand strategy decision trees for Ivan K and Sam 5 (Figs. 3.6–3.9).[10] Figure 3.6 shows how Ivan K might see the world at the time he begins an invasion of Southwest Asia *given* that he has already decided to

---

[9]See Jacobsen, Levine, and Schwabe (1985) for a discussion of the Europeans' planning from 1920–1940.

[10]Although these trees are much simplified for the purposes of a methodological paper, we note that in our experience even relatively simple trees capture much of what exists—i.e., grand strategy is often not very "grand." See also the literature on cognitive maps (Axelrod, 1976).

Attribute	Values of Attribute					
Use of Nuclear Weapons:	deterrent/ no-first-use	deterrent/ preemptive	deterrent/ escalative	warfighting/ no-first-use	warfighting/ no-first-use	warfighting/ escalative
Warfighting Style:	incremental/ responsive	incremental/ initiating	massive/ responsive	massive/ initiating		
Decision Style:	highly-analytic	analytic-mixed	reflexive			
Expansionism:	conservative	opportunistic	adventurous			
Perception of Opponent Commitment:	isolationist	vital-interests	realpolitik			

Attribute	Values of Attribute					
Use of Nuclear Weapons:	deterrent/ no-first-use	deterrent/ preemptive	deterrent/ escalative	warfighting/ no-first-use	warfighting/ preemptive	warfighting/ escalative
Warfighting Style:	incremental/ responsive	incremental/ initiating	massive/ responsive	massive/ initiating		
Decision Style:	highly-analytic	analytic-mixed	reflexive			
Commitment:	isolationist	vital-interests	realpolitik			
Alliance Dependence:	dependent	consultative	autonomous			

Fig. 3.3—A short list of attributes

Attribute	Values of Attribute					
Use of Nuclear Weapons:	deterrent/ no-first-use	deterrent/ preemptive	deterrent/ escalative	warfighting/ no-first-use	warfighting no-first-use	warfighting/ escalative
Warfighting Style:	incremental/ responsive	incremental/ initiating	massive/ responsive	massive/ initiating		
Decision Style:	highly-analytic	analytic-mixed	reflexive			
Expansionism:	conservative	opportunistic				
Perception of Opponent Commitment:	isolationist	vital-interests	realpolitik			

Fig. 3.4—Short list of attributes for Ivan K

Attribute	Values of Attribute					
Use of Nuclear Weapons:	deterrent/ no-first-use	deterrent/ preemptive	deterrent/ escalative	warfighting/ no-first-use	warfighting/ preemptive	warfighting/ escalative
Warfighting Style:	incremental/ responsive	incremental/ initiating	massive/ responsive	massive/ initiating		
Decision Style:	highly-analytic	analytic-mixed	reflexive			
Commitment:	isolationist	vital-interests	realpolitik			
Alliance Dependence:	dependent	consultative	autonomous			

Fig. 3.5—Short list of attributes for Sam 5

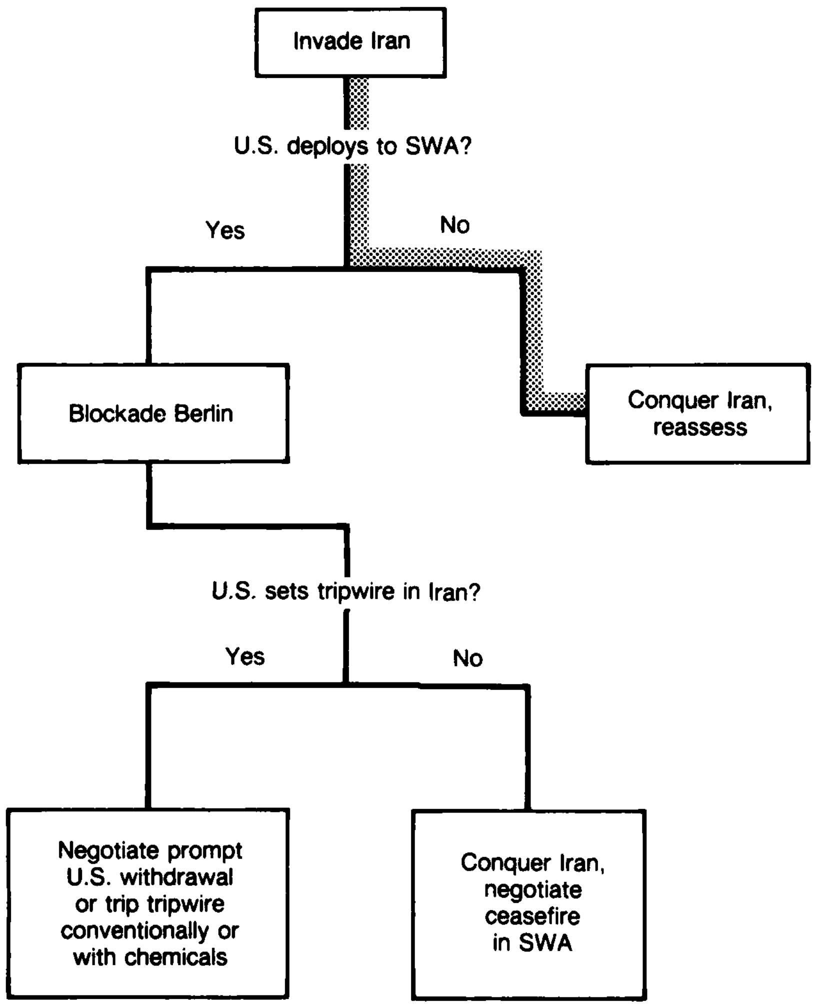

Fig. 3.6—Simplified grand strategy tree for Ivan K in SWA

take that step (for reasons of opportunism). Figure 3.7 shows how Ivan K might see the world later as he began an invasion of Western Europe (for reasons the tree alone does not show). Note that Ivan K would not invade Southwest Asia in the first place if he expected it to lead to war in Europe. Nonetheless, it is reasonable to ask how he would look at the world if he found himself in an unexpected situation. The shaded path indicates the path that Ivan K regards as most likely.

Without extensive discussion, let us merely note that Fig. 3.6 suggests that Ivan K invades Iran recognizing that the United States may deploy Central Command (i.e., rapid deployment) and other forces. He is prepared to make a feint in Berlin, or—if necessary—to trip the U.S. tripwire. He does not take all that seriously, however, the likelihood of a real European war.

Later, in Fig. 3.7, we see Ivan K taking an incremental view once again. The image he has is of a short conventional war, limited as much as possible (to the FRG and Denmark) so as to maximize the likelihood of making fundamental gains quickly and without nuclear weapons. On the other hand, he recognizes full well the possibility of a united NATO and of NATO first use of nuclear weapons. Thus, he will be prepared for those contingencies. He is relatively confident that war will not become intercontinental in scope.

Figures 3.8 and 3.9 describe the simplified views of Sam 5 for the same snapshots in time. The grand strategies are those discussed earlier in the prose essays.

Since it may not be immediately apparent why such decision trees are interesting, given their relationship to the prose descriptions and our insistence that they do not necessarily represent what the Ivans and Sams will in fact encounter or do in the simulation, we need to elaborate somewhat. The most significant feature of the grand strategy trees for Ivan K is that they show specific *incrementalist* strategies—something that may well be anathema to the Soviet military staff, but something that might be realistic for some Soviet leaderships nonetheless. In any case, Ivan K is an incrementalist at these levels of conflict. Should the war become nuclear, however, his war fighting style will be much more aggressive, decisive, and doctrinal. Other Ivans might never be incrementalist.

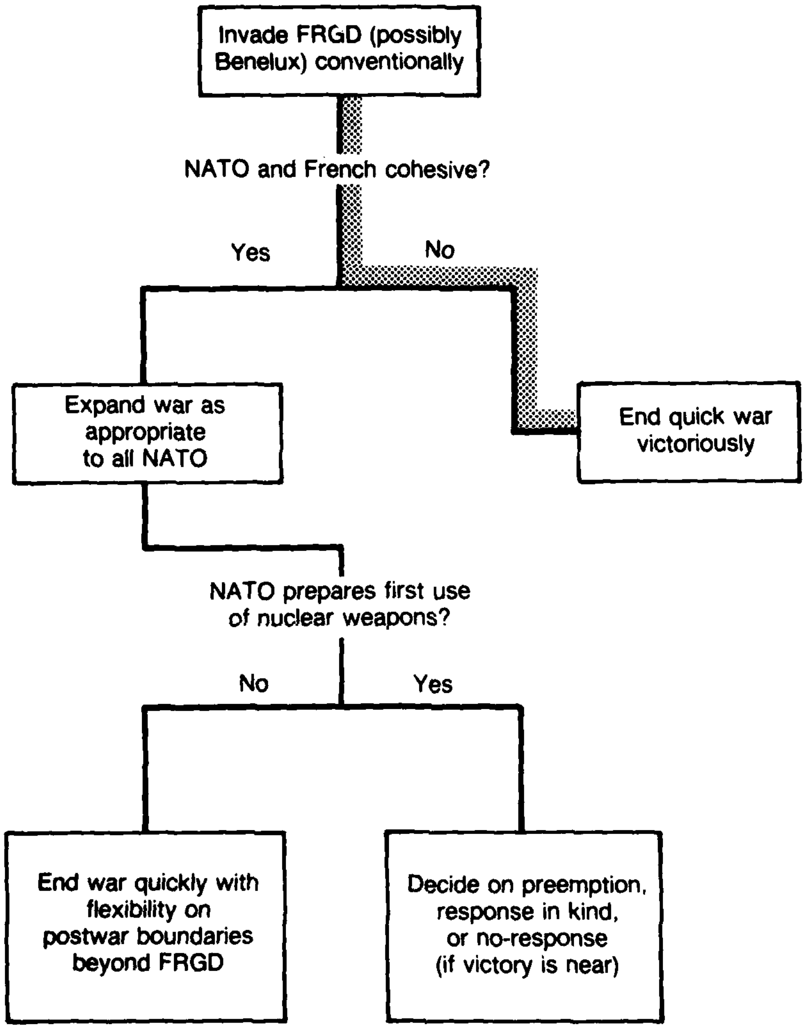

Fig. 3.7—Simplified grand strategy tree for Ivan K in Europe

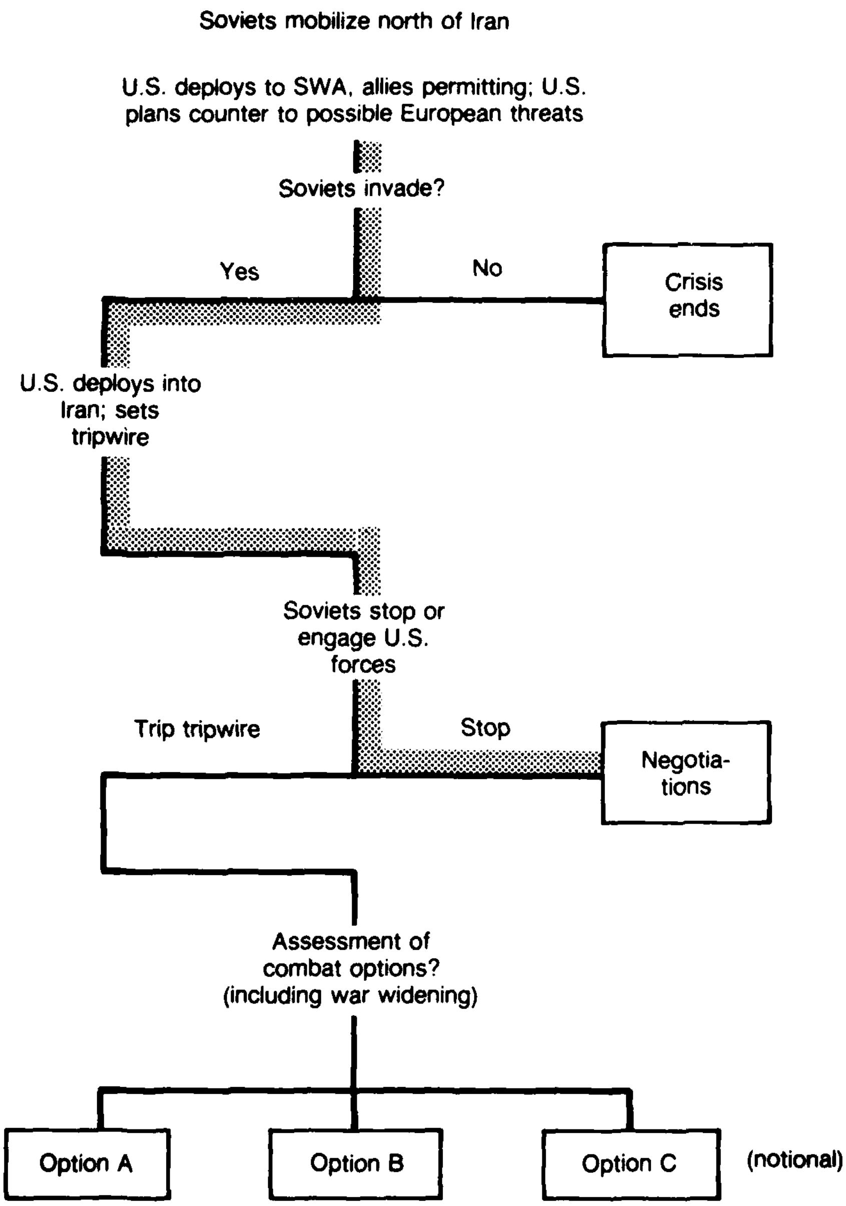

Fig. 3.8—Simplified grand strategy tree for Sam 5 in SWA

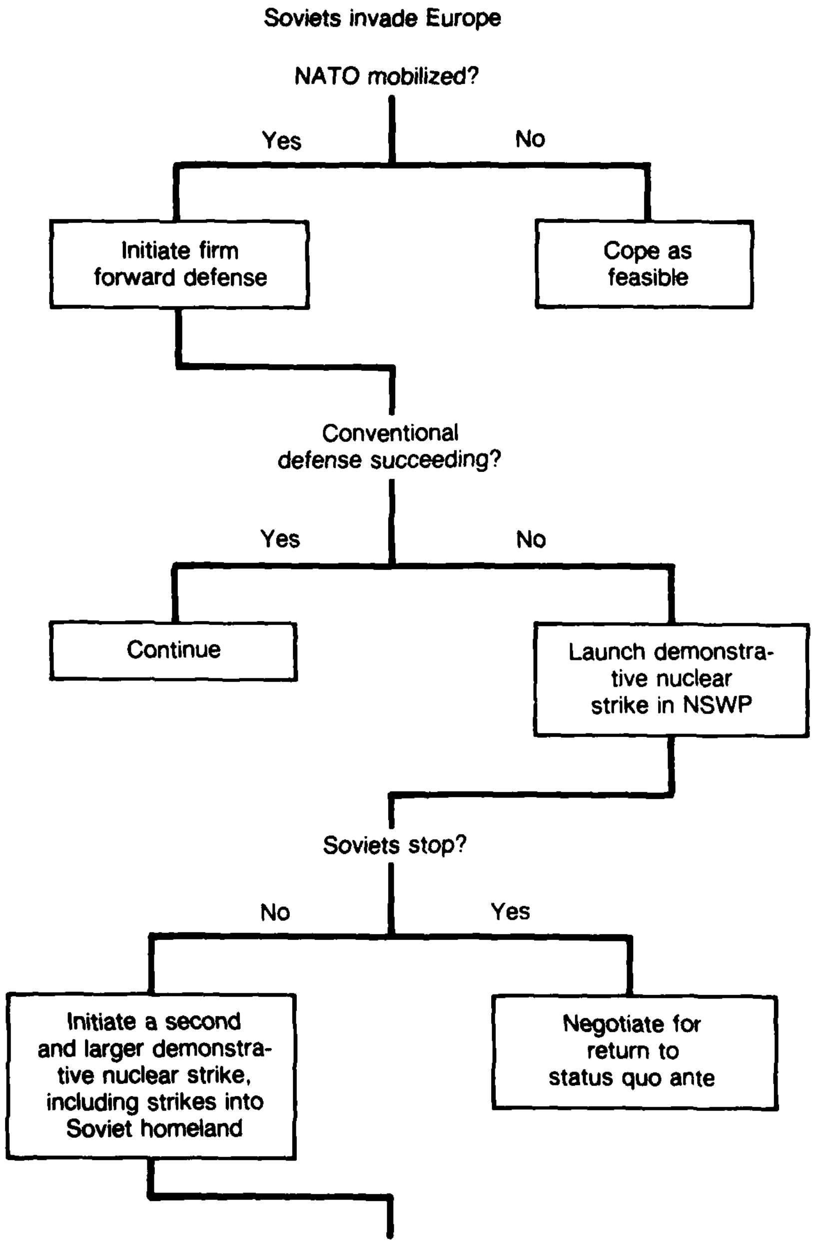

NOTE: Sam 5 does not really have a best estimate of outcome.

Fig. 3.9—Grand strategy tree for Sam 5 in Europe

A TRANSITION MATRIX FOR ESCALATION AND DEESCALATION

By the time the grand strategy decision tree has been constructed, the rule-writer should have a fairly clear idea of how his agent would behave in several circumstances. The next step, filling out a transition matrix for possible escalation and deescalation, is one more step in sharpening these ideas.

The 17 situations from Fig. 3.1 can be cast into a "from/to" matrix where the rows indicate situations that the agent might encounter and the columns indicate what new situation the agent might produce by his decision. This matrix we term an *escalation matrix*, or transition matrix, because it aids the rule-writer in considering what escalatory, deescalatory, or war spreading steps might seem particularly plausible and important for his particular Ivan or Sam. By concentrating on the most important cells of the matrix, the rule-writer discovers the transitions in escalation level for which rules are most urgently needed; other rules can evolve with time, thought, and experience using the model. Note that rules will eventually be needed for all situations, and that there are circumstances under which almost any Ivan or Sam should use nuclear weapons or terminate—regardless of what the rule-writer enters in the transition matrix as a guideline. The purpose of the transition matrix is merely to give the rule-writer a better handle on general propensities.

With this understood, then, one reads the matrices as follows. Given the limitation of problem scope decided upon in step 1, the bare matrix is the same for all Ivans and Sams; the rule-writer fills in the cell entries in creating his agent. As in the example matrices for Ivan K and Sam 5 shown in Figs. 3.10 and 3.11, an "o" entry indicates no escalation, a "+" entry marks escalation, and a "-" entry indicates deescalation. Blank cells are ones that the rule-writer can safely, at least at first, ignore.

GUIDING PRINCIPLES

By this point, the analyst should have a reasonably good sense of the principles that are guiding his judgments and should try to write them down. These guiding principles begin to look much more like rules.

ESCALATION GUIDANCE FOR IVAN K

Final State (Same numbering)

Initial state	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
1. Prep	o		+	+													
2. Term		o															
3. SWA-Sov-conv		–	o	+	+	+											
4. SWA-Sov-conv/Eur		–		o		+											
5. SWA-gen-conv		–			o	+	+										
6. Eur-demo-conv		–				o	+										
7. Eur-gen-conv		–					o				+						
8. Pnuc-Eur-gen-conv		–						o			+						
9. PIC-Eur-gen-conv		–							o		+						
10. Eur-demo-tac-nuc		–								o	+						
11. Eur-gen-tac-nuc		–									o			+		+	+
12. PIC-Eur-tac-nuc		–										o		+		+	+
13. Eur-demo-strat		–											o	+		+	+
14. Eur-CF-strat		–												o		+	+
15. Eur-strat		–													o	+	+
16. IC-CF-strat		–														o	+
17. IC-strat		–															o

Fig. 3.10—A transition matrix for Ivan K

ESCALATION GUIDANCE FOR SAM 5

Final State (Same numbering)

Initial state	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
1. Prep	o																
2. Term		o															
3. SWA-Sov-conv		–	o		+												
4. SWA-Sov-conv/Eur		–		o		+											
5. SWA-gen-conv		–			o												
6. Eur-demo-conv		–				o											
7. Eur-gen-conv		–					o			+							
8. Pnuc-Eur-gen-conv		–						o									
9. PIC-Eur-gen-conv		–							o								
10. Eur-demo-tac-nuc		–								o	+						
11. Eur-gen-tac-nuc		–									o		+				
12. PIC-Eur-tac-nuc		–										o					
13. Eur-demo-strat		–											o				
14. Eur-CF-strat		–												o			
15. Eur-strat		–													o		
16. IC-CF-strat		–														o	+
17. IC-strat		–															o

Fig. 3.11—A transition matrix for Sam 5

Ivan K's Guiding Principles

1. Termination is always acceptable if (a) the firm operational objectives have been met and Blue will not pose a serious threat; or (b) continuation would surely mean military defeat or horrific casualties while termination would not sacrifice the homeland.[11]

2. Termination is never acceptable if (a) Blue would continue fighting vigorously; or (b) there have been losses within the Pact and/or key occupied regions (FRGD and Northern Iran) and continuance would plausibly reverse them.

3. Escalation is acceptable: (a) in SWA; (b) in preemption; (c) in probable preemption in situations where the price of going second is high; (d) to improve results in a situation where termination is unacceptable; or (e) when escalating out of levels deemed to be inherently unstable or foolish (e.g., demonstrative use of tactical nuclear weapons). In these instances in which escalation is acceptable the decision would be made on the basis of technical considerations—i.e., on balance, whether the escalation would be expected to benefit the military effort.

4. Deescalation from intercontinental to theater strategic, or from theater-nuclear to postnuclear conventional, would be unacceptable except for tactical reasons. Should there be a phase in which nuclear weapons were not used somewhere, Ivan K would use them again whenever he deemed it tactically appropriate.

Sam 5's Guiding Principles Decision Tree

In some instances it is convenient to express principles as another type of tree such as that shown below for Sam 5 (Fig. 3.12). Whether one prefers verbal principles or trees is a matter of taste and situation. In any case, the principles upon which Sam 5 governs his action are fourfold. First, Sam 5 has a flexible response to aggression. He will deter if at all possible, will fight aggression when deterrence fails, and will escalate (both horizontally and vertically) when necessary to achieve these objectives, in incremental fashion. Second, Sam will bargain when he has a position of strength and will avoid bargaining from weakness if at all possible. Third, Sam 5 will avoid risks unless the stakes are high; i.e., unless the United States or a vital interest is

---

[11]It can be argued that Ivan K "should" be willing to terminate if prospects for success are virtually zero, even if continuance would not mean defeat or horrific casualties. However, it has been notoriously difficult for nations to terminate conflict on the basis of such rational calculations in the past and we assume that Ivan K would simply not back away from his objectives unless forced to (at least for the relatively short wars we are currently simulating).

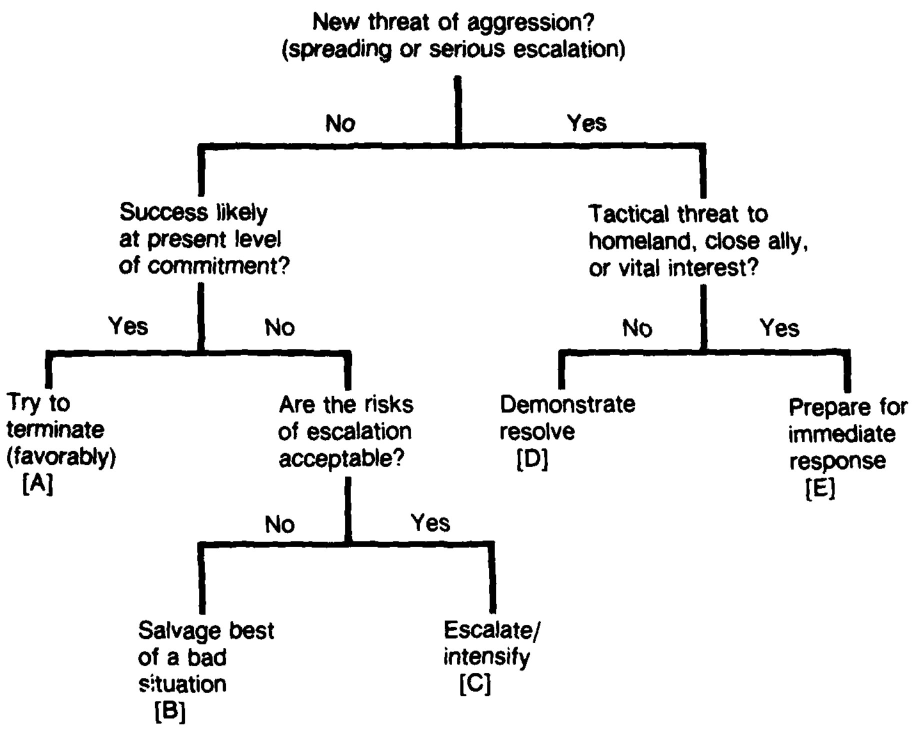

Fig. 3.12—A guiding principles decision tree for Sam 5

threatened. Finally, Sam will seek allied support and will share responsibility as much as possible, and particularly in cases when the war remains centered in nonhomeland theaters. Each of these principles arises from the verbal description as carried through the subsequent steps of agent construction and is expressed in the tree shown in Fig. 3.12.

SUMMARY OF CHARACTERIZATION TECHNIQUES

At this point, finally, an analyst should be ready to actually develop coherent rule-sets for his Ivan or Sam. He will have several different packages indicating the agent's predilections in one form or another (essay, grand strategy trees, and transition matrices), he will have a more rigorous set of guiding principles by which he will try to live in writing rules, and he has a set of formal attributes. Depending on his patience and interest in preliminary work, he may have iterated all these items to make them more nearly consistent. Thus, his essay may have taken on some of the content of his grand strategy trees and his transition matrix may have gained so many x's and o's as to be less useful than in the first run-through when he was defining predilections. In any case, however, he should now be ready to proceed. Next, he wants to embody his model in a computerized framework. That is the subject of Secs. IV and V.

IV. AGGREGATED SYSTEM SPECIFICATIONS FOR THE NCL MODEL

OVERVIEW

This section and the next describe the transition from a conceptual NCL model as sketched out in Sec. II to the prototype computer program we are now using and building upon. This section describes the *technical requirements* for the computer program (i.e., it describes the "system specifications"). It therefore deals with issues such as the user interface, inputs, outputs, and the relationship architecturally between the NCL models and the overall game-structured simulation. Figure 4.1 itemizes those issues. Because the section is inherently technical, some readers may wish merely to skim it.

- Modes of operation
- User requirements
- Requirements imposed by architecture of overall RSAC simulation:
  - — Object classes
  - — Control flow among objects
  - — Data flow among objects
  - — Interface architecture, including inputs and outputs
  - — Look-ahead mechanisms with treatment of perceptions

Fig. 4.1—Issues for basic system specification

MODES OF OPERATION

One of the most fundamental issues in system specification is recognizing the modes in which the model is to operate. **We have required four distinct modes:**

- *Integrated operations* in which the NCL models operate within the RSAC's overall game-structured simulation.
- *Support operations* in which the NCL models act as decision aids to provide information and advice to human teams charged with NCL decisions.
- *Standalone operations* in which the NCL model interacts only with an analyst (i.e., with no use of the other models).
- *Partial-system operations*, most notably "flag wars" in which two competing NCL models and possibly the Green Agent interact without operation of the Force Agent by signaling their actions to one another through the posting of symbolic messages ("flags") in the world state.

The first two modes are probably obvious, the others less so. The standalone mode is important because it allows analysts focused on NCL issues to interact exclusively with the NCL model without always being burdened with the complexity of full-system operations. The partial-system mode allows the analyst to see action-reaction phenomena at the political level without having to operate the Force Agent's simulation. Having these modes is also useful in debugging and maintaining software, and in hedging against problems in one part of the overall simulation bringing all substantive work to a standstill.

These requirements for operational diversity have a significant effect on both interfaces and control logic, as will be clearer in the context of a particular program implementation (see Sec. V). The key issue is having unobtrusive "boilerplate code." By boilerplate code we mean code written by computer programmers to provide invisible services such as telling the computer whether to expect inputs typed by the user (and how to process them), or whether to expect inputs from fixed data bases (and where to find them in the data bases). The user should not have to do anything more complicated than, for example, choosing "human mode" versus "automated mode" from a menu; after that, the computer should know what to expect and do. **At the level of system specification, the requirement is that boilerplate code related to operational mode be as separate as possible from the substantive code of concern to the analyst developing NCL decision rules.**

THE USER INTERFACE: ACHIEVING TRANSPARENCY

Methods for Achieving Transparency

Most computer models are used routinely without looking into their internal logic except in a general way. For example, a user may read in a model's documentation that at one point the model calculates an object's velocity using $V(t) = Vo + \int F(s)/M\, ds$. He recognizes this as an expression of Newton's laws and reads on, without worrying about how the computer program evaluates the integral. Instead, he worries about what he has to specify as input data (Vo, the function F(t), and M).

By contrast, the NCL models are of a very different nature. The models are not based on empirical laws of nature but on a collection of hypothesized decision rules reflecting a myriad of judgments. Even if the models are built by intelligent analysts, neither we nor our sponsors have interest in a "black box" NCL that would make mysterious and arbitrary decisions in war games or simulations regarding such cosmic issues as escalation, termination, or choice of military strategy. There are *fundamental* uncertainties about real-world decisionmaking on such matters, and there is no way even in principle to validate an NCL model's predictive performance. Thus, we are concerned not only with the surface plausibility and apparent coherence of the model's performance, but also with the model's logic (and the implications of alternative logic). This implies that we must require a high degree of model transparency.

Some rule-based programs achieve a degree of transparency by allowing the user to express the rules as data, without asking the user to look at computer code. However, this does not avoid the problem of transparency; when the collection of rules becomes quite large, understanding their effect can be just as difficult as reading computer code. Those using NCL models are interested in the inference mechanism as well as the individual decision rules, are interested in changing the rules in diverse and complex ways, and are working with models that will continue to change. All of this implied to us the need to make the underlying computer code as transparent as possible.

To achieve transparency at the level of the computer code we have required and developed the following:

- A computer language that can be read and understood by people with subject-area knowledge but only minimal knowledge of computers;
- A separation of substantive and boilerplate code in the NCL programs;

- "Structured programming" with the substantive content of the code organized precisely along the lines of the conceptual model described in Sec. II;
- Software facilities to assist the analyst in following the model's logic, notably a data editor and an automatically generated explanation log (other aids are also under development); and
- A computer language that executes quickly enough to permit the man-machine interactions needed for full comprehension of complex simulations.[1]

Requiring this level of transparency is unusual. Even in academic artificial intelligence research, few projects attempt to provide all of the above features. However, **our intention is that analysts and other substantive experts work directly with the software tools to build, test, use, and adjust the models.** We believe that computer technology will now permit true man-machine cooperation greatly extending the scope and complexity of what an analyst can comprehend. Thus, we intend to minimize intermediaries such as technician programmers and unfriendly programming languages. Such intermediaries have the effect of breaking concentration and interfering with creative thinking. They also create the peril of divergence between the modeler's intent and the code's reality. Such divergences are notoriously hard to detect, and the mere possibility of such divergences can greatly reduce the perceived value of a computer program.

The Rand-Abel™ Programming Language

Perhaps the most important element of the user interface in a rule-based model that is to be modified frequently is the programming language, in our case Rand-Abel, which was developed specifically for building models such as those discussed in this report (Shapiro, Hall, Anderson, and LaCasse, 1985a, 1985b). The Rand-Abel programming language's principal characteristics include:

- *Clarity* (key segments of Rand-Abel code can be read by nonprogrammers so long as they understand the subject matter, including its jargon);

[1]Influenced by our colleague Norman Shapiro on this matter, we seek execution speeds of *seconds* so that the analyst can interact with the model within his short-term cognitive cycle—i.e., the period over which the analyst is uninterruptedly focused and efficient. With complex simulations it is necessary, not merely desirable, for the analyst to "work with the model" extensively if he is to understand it (he must also sit back and *think* about the models' underlying relationships, but that is another subject).

- *Speed* (within a factor of three of the speed of the C programming language; in our context, less than one millisecond per rule on an unloaded VAX 11-780);
- The use directly in code of *decision tables* essentially identical to what an analyst might work out on the back of an envelope (such tables are closely related to decision trees); and
- Very *strong typing* to catch major sources of traditional software errors early.[2]

Two examples of Rand-Abel code follow. The first gives a decision rule in prose form; the second uses the decision table mechanism mentioned above. Again let us emphasize that what follows is executable computer code.

```
If Current-situation is SWA-Sov-conv
Then
{
 If RDF-deploying is Yes
 Then [from Red's point of view]
    Let Strategic-warning be SWA-conv.
    Else    Let Strategic-warning be none.
}
```

Before discussing this Ivan K rule, we should mention two conventions in Rand-Abel. First, hyphenated words are really single variables. Second, anything between brackets [ ] is a comment ignored by the computer. This rule says that if the Soviets (but not the United States) are engaged in conventional conflict in Southwest Asia (a conflict level we call SWA-Sov-conv), then they will consider deployment of the U.S. Rapid Deployment Force as strategic warning of war with the United States in Southwest Asia (a level of war we call SWA-conv). This is hardly profound but providing large numbers of such simple rules is essential in giving intelligence to the computerized NCL model. In passing, note that the rule is readable if and only if one

[2]"Strong typing" is a mechanism for preventing a large class of errors in which the user expresses a relationship between two variables that appears plausible to the computer but which is actually nonsensical. When introducing a new variable, the user must specify its "type" (e.g., integer or qualitative variable). If it is a qualitative variable, he must specify the range of permitted values (e.g. the alert state may be low, medium, or high). In Rand-Abel, only variables with *identical* ranges are of the same type and only variables of the same type can be related to one another in expressions. This enforces a kind of modularity and rigor in rule-writing that compensates for the fact that many English words used in rules have different meanings in different contexts.

knows the jargon of the subject area because, in practice, one wants to use abbreviations in rule-writing.[3]

The second example is a truncated and simplified version of a notional decision table for Sam 5 fighting a conventional European war (i.e., the current situation to which the table applies is Eur-gen-conv).

Decision Table

Main-theater-status	Other-status	Prospects	Risks	/ Escalation-guidance
goals-met	good	--	--	Terminate
[Many lines omitted for brevity]				
progress	good	good	low	Eur-gen-conv
no-progress	bad	bad	--	Eur-demo-nuc.

To understand the decision table, let us consider the first line below the table's header line defining the columns. The computer will read this line to mean "If Main-theater status is goals-met and Other-status is good, then regardless of Prospects or Risks, Terminate." That is, the headers are the names of variables, with those to the left of the "/" being independent variables, and that to the right being the dependent variable (i.e., the decision). The symbol "--" means that this particular rule does not depend on the variable in the column in which the "--" appears.

Our decision tables can be equivalent to decision trees, with one line for each end point of a tree. Thus, not only can we express rules easily that were developed with decision-tree logic, we can also be confident that we have *all* the cases covered in a set of rules (a traditional problem with more conventional If . . . Then . . . Else statements). In practice, we usually condense the tables by using the "--" feature, and by mid-1986 we shall be able in tables to have entries such as ">low" meaning, for example, that the variable value is "medium or high" if the variable has possible values of low, medium, and high. Use of >, <, and similar operators is already permitted by Rand-Abel in ordinary (nontable) statements.

[3]Although there are syntax restrictions, we *could* write Rand-Abel code in long and flowery sentences with only a few strange symbols and conventions. Thus, abbreviations like SWA-conv and SWA-Sov-conv are purely for analyst convenience and not a feature of the language.

Our experience with Rand-Abel has been quite positive. Although we are now adding features to the language, notably interpretive execution and sets as data types, it is already practical for analysts with modest programming skills to read the substantive portions of NCL code and modify it themselves. It is also possible for analysts to develop all of the substantive rules in quasi Rand-Abel, and then to review the final products developed by professional programmers.

REQUIREMENTS IMPOSED BY THE OVERALL RSAC SIMULATION SYSTEM

Obviously, the NCL model must fit within the overall game-structured simulation described conceptually in earlier sections. It is therefore useful to describe that simulation more technically before proceeding.

The Analyst's Conceptual Model

First, let us review the RSAC simulation as it might be characterized by the strategic analyst concerned with substantive *modeling*. Figures 4.2–4.5 are in this character. Figure 4.2 is the simplest image of all, establishing only that the simulation is more or less a two-sided game with the environment of the players (Red and Blue) modulated adaptively by the Control Agents (Green and Force).[4]

Figure 4.3 provides a more complex image, the essence of which was developed in 1982–1983 and used as a top-level conceptual model driving system requirements (Davis and Stan, 1984; Davis, 1984). In this conceptual model:

- The simulation proceeds (first column) with the various agents taking turns, although not in a particular order, and not with perfect information about what the other agents have recently done. The agents' turns are determined by wakeup rules, which *they* specify on the basis of either events or elapsed time and a protocol for deciding which moves first when more than one agent wants to move at the same time.

---

[4]The figure, which dates back to 1980, is somewhat misleading in that the simulation is technically an n-player game, since third countries make independent *political* decisions on an equal basis with Red and Blue (although with less sophisticated logic). See Schwabe and Jamison (1982) and Shlapak, Schwabe, Ben-Horin, and Lorell (1985). Alternatively, the simulation can be regarded as a two-player game with additional "dummies" (see Shubik, 1982).

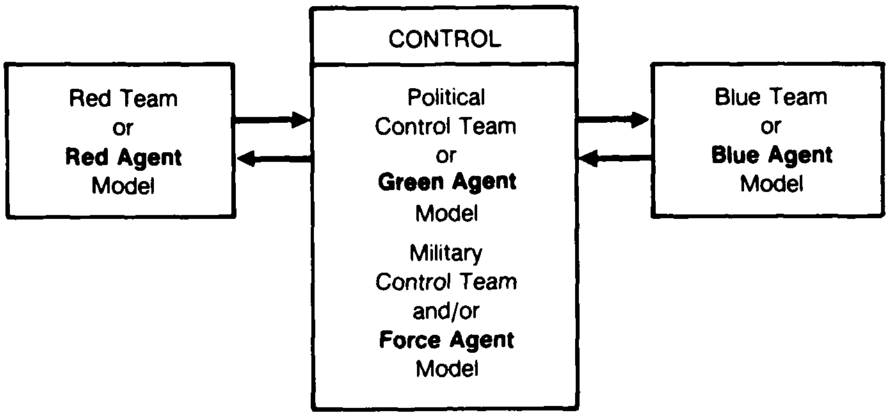

Red Agent: Makes decisions for Soviet Union and Warsaw Pact
Blue Agent: Makes decisions for United States and NATO
Green Agent: Makes decisions for nonsuperpower countries
Force Agent: Simulates results of combat and other military operations; also, controls simulation time

Fig. 4.2—Simplest image of the RSAC simulation

- Red and Blue operate on the basis of *analytic war plans*, which provide instructions (including adaptation rules) for both military and diplomatic actions worldwide. Both Red and Blue have a set of analytic war plans available for use in the particular game. These must have the character of building blocks to permit changing plans in the midst of a simulation and to separate actions by theater, force type, campaign phase, etc. They are *not* hard-wired decision trees, because we cannot even predict the order of events with confidence, much less whether they will all occur, or when.
- When Red takes a turn (second column), it first decides whether the current analytic war plan is still appropriate. If it is, then it continues to use that plan, which is implemented by the military command levels in the form of orders and messages to the Force Agent, Green Agent, and Blue Agent. If the plan is no longer appropriate (on the basis of rules contained within

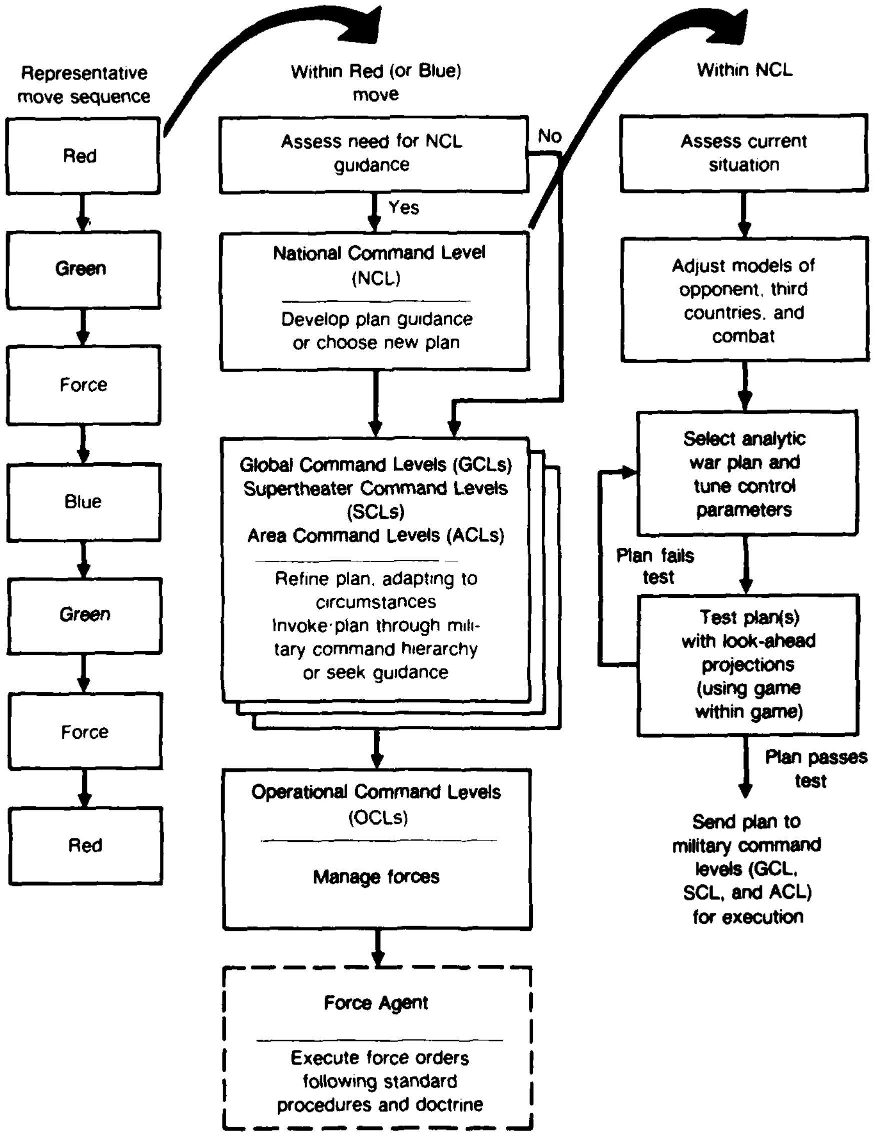

Fig. 4.3—Conceptual model of the RSAC simulation

that plan—rules that can recognize opportunities, failures, and mission completion), then the NCL model is activated to make decisions about changing the analytic war plan.

- The NCL goes through a decision process indicated in the third column—a process that includes situation assessment, learning, and choice of a new analytic war plan (the process for which will be discussed in Sec. V).
- In deciding on a new analytic war plan, the NCL may use one or more look-ahead projections in which the full simulation is run within itself (a game within a game), but with other agents replaced by Red's perceptions of reality (Red's Blue, Red's Force, and Red's Green). If the tentative plan fails in look-ahead, the NCL will choose a new one—hence, the feedback shown in column three.
- The actual execution of force orders, along with tactical-level decisionmaking as needed, takes place in the Force Agent rather than the Red and Blue Agent (the Force Agent can use different models for Red and Blue tactical-level decision processes).

Figure 4.4 is a simplified view of the Red Agent structure showing that the Red (and Blue) Agent is actually a hierarchy of models corresponding to different real-world commands. The Global Command Level corresponds roughly to the U.S. Joint Chiefs of Staff or the Soviet Supreme High Command; the Supertheater Command(s), of which there may be none or several, correspond to commanders like the U.S. CINCEUR or the NATO SACEUR, who have multitheater responsibility, and the Area Commands correspond to U.S. CINCs or Soviet TVD commanders. Not shown here are the operational and tactical command levels, which are modeled differently.

Finally, Fig. 4.5 is an *influence diagram* showing how the various agents and agent components should interact functionally. We require a line of authority by which the NCL communicates directly with the GCL, but not with the SCLs.[5] Only the GCL exchanges messages with third countries (i.e., with Green Agent).

This, then, summarizes the basic conceptual model of RSAC system operations within which the NCL model is to operate. These representations are essentially those of an analyst rather than a software

[5]This system specification does *not* limit in a practical sense our ability to model the skipping of echelons in command-control. Whether an intermediate echelon of command worries about an issue or passes it on without delay or amendment is a matter for the substantive rules.

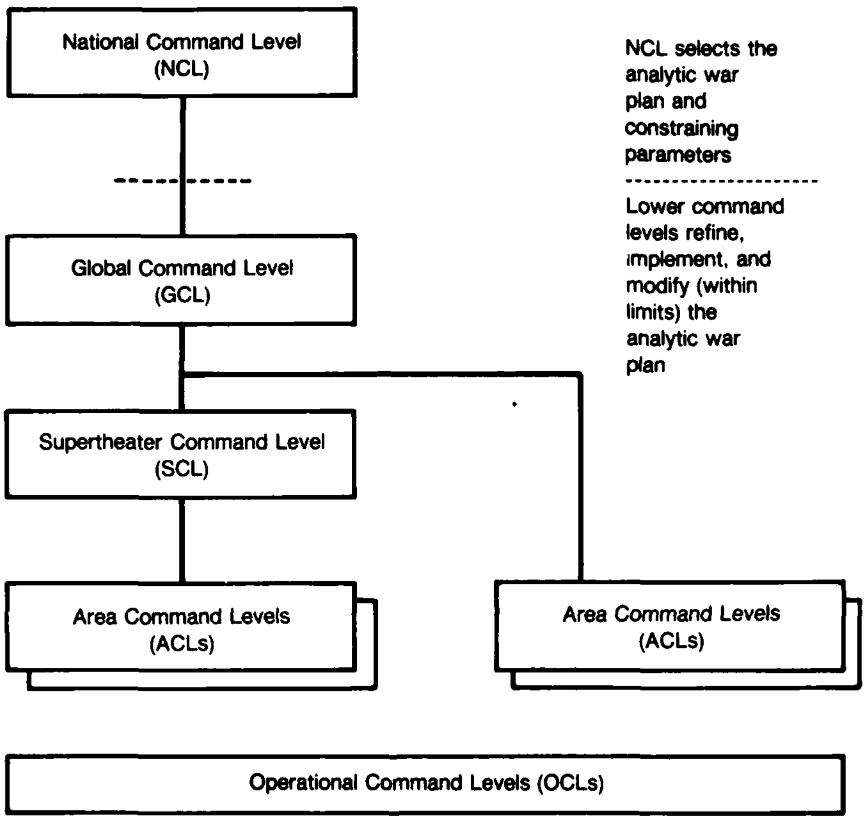

Fig. 4.4—Simplified view of command levels within the Red and Blue Agent (Red may not have SCLs in practice)

designer. They have major implications, however, for the effective flow of both control and data in a working system.

Specifications for System Software

Turning now to the system-software specifications, Fig. 4.6 shows the subprograms (objects) of the simulation.[6] System Monitor is the main program and, at least figuratively (see below), "calls" the others

[6]The Control Agent shown on the left side is a new (Spring, 1986) mechanism allowing the analyst to schedule certain events to occur (on the basis of time or world state) that he would otherwise have to introduce interactively. The Control Agent has rules and turns but is not a "player."

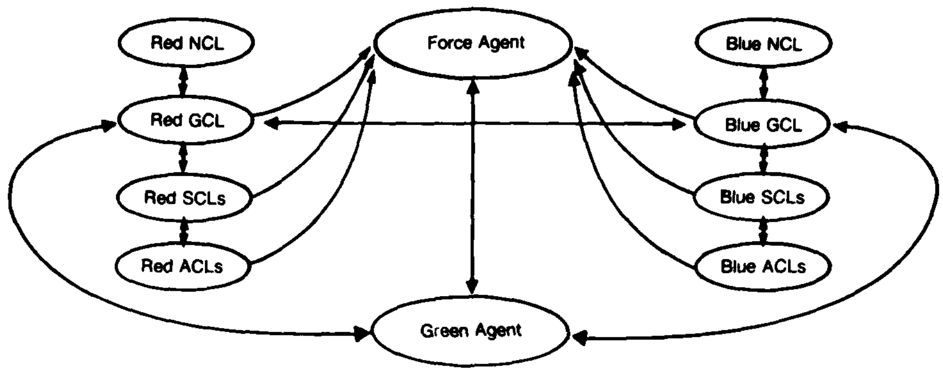

Not shown: All models read world-state data reflecting actions of Force Agent and Green Agent.

Fig. 4.5—Influence diagram for the RSAC simulation

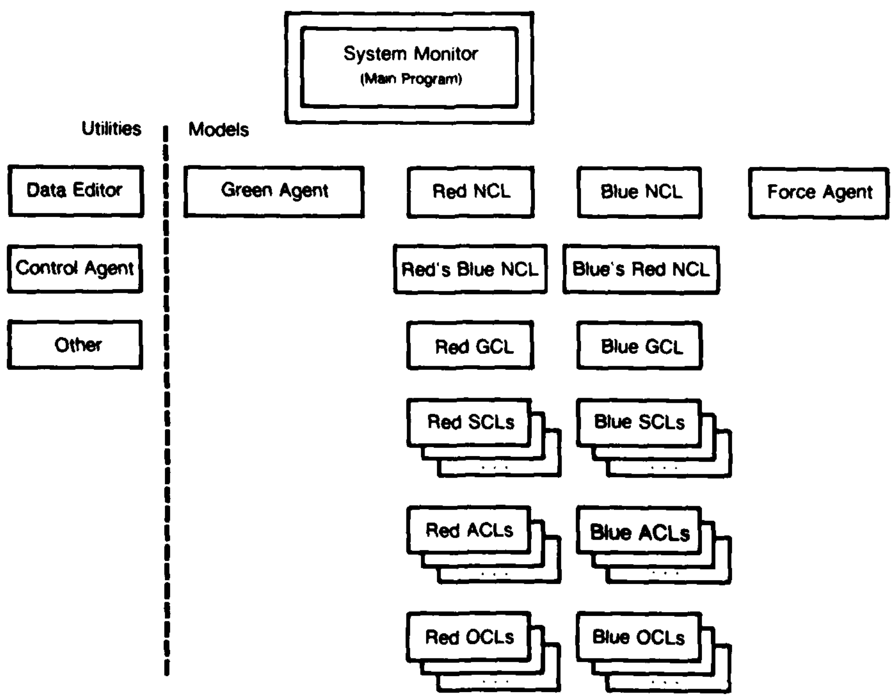

- All subprograms can be separate coprocesses
- "Red Agent" and "Blue Agent" are abstractions

Fig. 4.6—Top-level subprograms in the RSAC simulation

as subordinate programs—i.e., System Monitor is a software entity that determines which of the other agents (or human teams) have control. We cannot show a control flow chart because the order in which the various agents move depends on events in the simulation. We also do not indicate the *number* or identity of the various SCL and ACL models because those depend on the particular strategic framework being used in a study.[7]

Perhaps the most important technical requirement here stems from our desire to have many of the objects of Fig. 4.6 display the character of organizations (military commands) following script-like orders without constantly having to reassess the full range of global issues. If the computer programs representing these objects are to convey that image,[8] it is necessary that they be written so that when, for example, a Red ACL commander has a turn for decisionmaking, the computer merely goes to the place in that Red ACL's plan where decisions were made last rather than starting at the top of some general subroutine program. The metaphor used here is as follows: A given submodel such as the Red ACL in question makes its decision and then "sleeps." When next it "awakens," it begins where it left off and continues down the plan.

Although other implementing mechanisms are possible in principle, **this particular system specification virtually dictates the need to represent the objects of our simulation as coprocesses** (Knuth, 1973:190)[9] rather than as subroutines.[10]

One subtlety in Fig. 4.6 is the appearance of Red's Blue NCL and Blue's Red NCL. As discussed above, some decisions by Red and Blue depend on their *perceptions* of reality. Although some of these perceptions can be represented by changes in *parameters* (i.e., parameters affecting particular Force Agent calculations or the modeling of third-country decisions in Green Agent), others require separate models alto-

---

[7]Some RSAC publications (e.g., Shukiar, 1985) highlight *wakeup rules* as a separate module for pedagogical purposes. However, from an architectural viewpoint, the wakeup rules reside in functions within the Red and Blue Agents. The Force Agent conducting the simulation calls on those functions to determine when agents wish to move; when it detects one of those wakeup conditions it passes control to System Monitor, which then determines what should be awakened first and passes control accordingly.

[8]Having the program structures mirror the conceptual model's structures is important for both transparency and for allowing analysts of modest programming skill to work interactively with the computer in adapting the programs to their desires.

[9]Most of the objects in Fig. 4.6 are currently implemented as separate coprocesses. The exceptions are that Force Agent, Data Editor, and System Monitor share a coprocess, as do Blue's Red and Red NCLs and Red's Blue and Blue NCLs. Objects sharing a coprocess are subroutines within the same coprocess.

[10]Implementing the concept of hierarchical coprocesses within the environment of Rand-Abel, C, and Unix was accomplished by Rand colleague Ed Hall.

gether. The opponent-NCL models are the principal example, but the architecture also permits us to have perceived opponent war plans.

Figure 4.7 shows a simplified and idealized representation of the system's required data flow. It is adequate for most purposes but not for establishing interface details. It omits System Monitor, to which all the objects of Fig. 4.7 are connected.

Figure 4.8 shows a more detailed representation of system specifications for the data flow within the RSAC simulation, one showing the individual interface programs and major data sets. Again, however, we omit flows to System Monitor. Each model within the decision-model box is independent of the others. Moreover, the Red and Blue Agents are each composed of many technically independent submodels as indicated in Fig. 4.6. **Notable features of this specification include:**

- The system's natural "objects" (i.e., Green Agent, Force Agent, and the various independent components of the Red and Blue

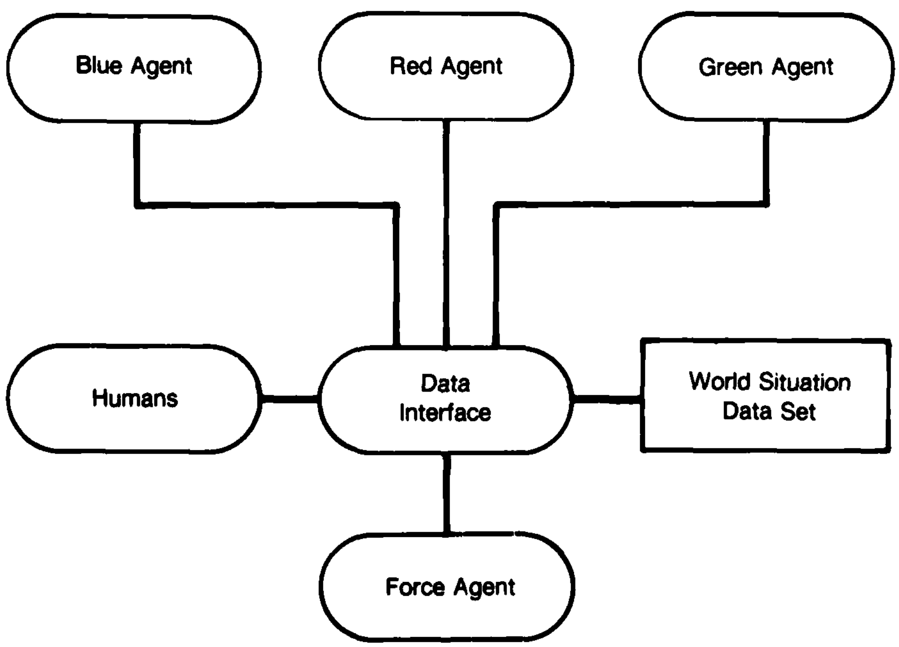

Fig. 4.7—An idealized representation of data flow

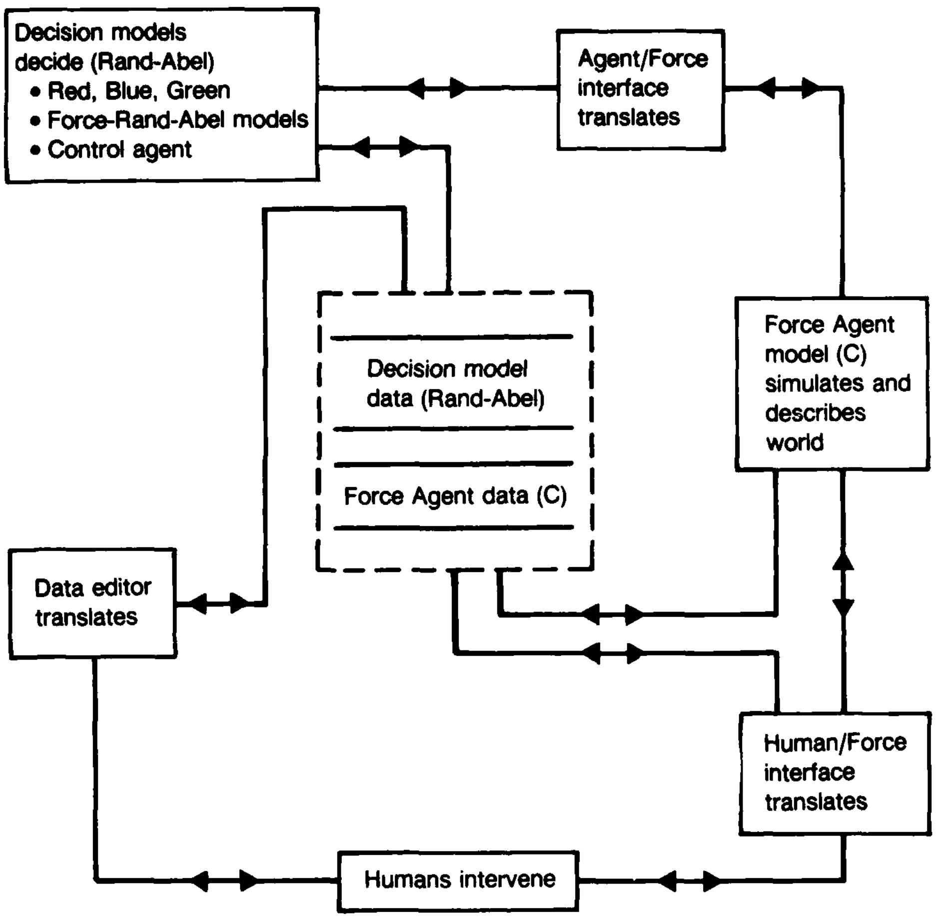

Fig. 4.8—System specifications for data flow

Agent) communicate technically by sending messages.[11] These messages are stored in the World Situation Data Set (WSDS), allowing the models to share information while still providing strong controls over what data a model can access.

[11]The RSAC simulation has many of the same features as simulations based on so-called "object oriented programming languages" such as Smalltalk (Goldberg and Kay, 1976) and ROSS (McArthur and Klahr, 1982), but the relationships are subtle and the mechanisms for achieving the same functionality quite different. Although the literature is inconsistent, most object oriented languages include *message passing*, *data hiding*, and

- The system divides into a rule-based part written in Rand-Abel (Red Agent, Blue Agent, Green Agent, some special Force Agent rule-based models, and some rule-based models that allow the analyst to schedule interventions much as would a control team) and a portion written in the C language. The latter is concerned with traditional quantitative simulation of physical events, the keeping of the clock, and so forth.
- Interface programs exist to control the flow of information,[12] as well as to translate information and, in some instances, to operate on data to produce useful aggregations.
- The sum of all information describing the state of the game is contained in a WSDS, which is manipulated by the System Monitor in performing look-aheads. The WSDS has two independent parts, one serving the Rand-Abel models and one serving the C models.
- Humans can interact directly with the Force Agent by issuing commands. They can also obtain information directly from the Force Agent (via an interface) and can change parameters in the simulation. They can interact with the decision models through the Data Editor.[13]
- In conducting look-aheads, the system uses a copy of the complete WSDS, the original WSDS having been pushed onto a stack that allows arbitrarily many look-ahead branchpoints. It is possible to rewind the game back to any such branchpoint, thus making it easy to explore different branches of the scenario space without starting over again.[14]

---

inheritance; interactions among objects and the environment occur only through exchange of messages—i.e., they occur only at discrete times. The RSAC simulation has a form of message passing and substantial data hiding (implemented through Rand-Abel's strong typing features and the coprocess representation, which retains memory and context), but no inheritance. Interactions can be discrete or essentially continuous (e.g., by an ultimatum or nuclear attack on the one extreme, or by a continuing conventional war of attrition on the other). Most important, however, both the current work and object oriented modeling (as well as much traditional simulation with SIMSCRIPT) attempt to use representations that are physically natural to the problem.

[12]The mechanism for denying the Red Agent knowledge of the Blue Agent's internal thinking is the Rand-Abel language's very strong typing. Variables are typed in part by who "owns" them, and variables owned by different subprograms are different variables. Attempts in one subprogram to use a variable defined in another subprogram will lead to an error message if that communication is prohibited. Enforcement of such matters is possible because all Rand-Abel communications occur through a data dictionary.

[13]Since early 1986 the two interfaces with Force models have been virtually identical technically. For example, the mechanism by which a decision model obtains the value of a particular Force Agent variable is precisely the same as the mechanism by which a human player would be provided that information in a display. Also, any command or parameter setting that a human player can issue directly to the Force Agent can be issued instead by a decision model using the same syntax.

[14]Since "rewind" capabilities are often discussed in the abstract, we should note that this and nearly all the other aspects of the system described here have been operational in the RSAC program since 1984.

INTERFACES: INPUTS AND OUTPUTS

General Comments

With this background of system specifications for the overall RSAC simulation, we can now discuss inputs and outputs, both substantively and in terms of interface requirements. Figure 4.9 shows one view of the issue, distinguishing between final and intermediate outputs but remaining at a relatively high level of generality. **Figure 4.10 shows inputs and final outputs in more detail.** These follow from the conceptual model discussed earlier, the overall system architecture, and (in the case of the "flags" shown as the last type of output), from the need for certain special types of operational mode. We shall discuss that in more detail below.

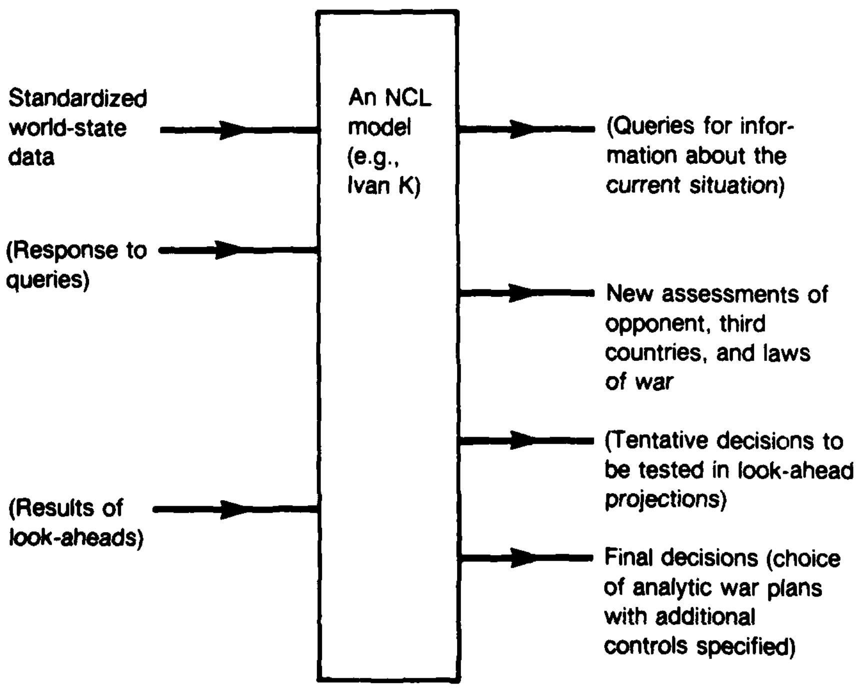

Convention: parentheses indicate data flow *during* NCL decision process, not final inputs or outputs.

Fig. 4.9—Generic input-output relationships for the NCL models

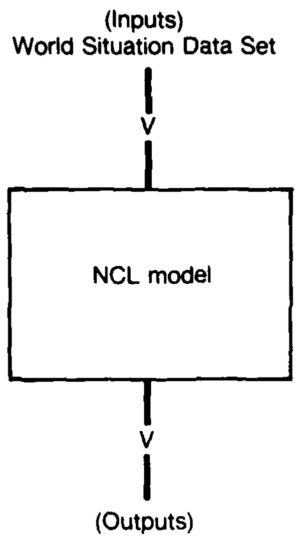

- Plan names, by command
- Values of plans' control variables (defaults exist)
- Name of opponent model
- Parameter values for third-country perceived behavior
- Parameter values for models of force operations and combat used by the agent (his perceptions)
- Variable values specified directly to opponent NCL, if any ("flags")

Fig. 4.10—NCL model's inputs and outputs

Inputs to NCL Decisions

As indicated above, the NCL communicates with *all* other objects in the RSAC simulation by way of the WSDS and the data interface. In a sense, then, there is only one interface. However, because the data interface performs filtering and other operations, in part to assure that Red does not know everything Blue knows and so forth, and in part because there are two types of data base within the WSDS (see Fig.

4.8), the reality is that we must specify the various interfaces separately.

Thinking of the NCL's inputs as information, **overall system design requires that we distinguish two types of input interface**:

- A ("direct-input approach" in which the NCL obtains information directly from the decision model (Rand-Abel) data base; and
- One in which the NCL obtains information indirectly from the Force Agent's data base (C) by using a special *query function* (also called an access function).

Implementation of these requirements is largely a matter of the programming language rather than the NCL program and can therefore be discussed here rather than in the next section.

An example of the direct-input mode would be the Rand-Abel statement seen earlier:

If RDF-deploying is Yes Then . . .

where RDF-deploying is a variable in the Rand-Abel data base. The value of RDF-deploying could be set in one part of the model and used elsewhere, could be set on a previous wakeup and used later, or, if set by the analyst, could be viewed as an input to the model.

The following is an example in Rand-Abel of the query-function approach:

```
Let Eur-Blue-attrition be the report from Ask-force-count using
        Troops as unit-type,
        All as unit-owner,
        Blue as side,
        All as in-area,
        Normal as minimum-status,
        Normal as weapons-level,
        All as region-owner,
        CEur as assigned-to, and
        Cumulative-attrition as what-to-count.
```

Here Ask-force-count is a query function using Force data as input. Obviously, this type of Rand-Abel statement is *not* transparent unless one knows the syntax of the Ask-Force-count function. Thus, we attempt to hide such statements in service code that analysts need not see and to have analysts work almost exclusively in the direct-input mode.

Standard Outputs

Figure 4.10 defined the NCL model's outputs, but it may be useful to give some examples and elaboration.

1. *Plan names.* In saying that the NCL *selects an* analytic war plan, we should emphasize that the plan has many components—one for each of the objects in the hierarchy shown in Fig. 4.10. Thus, the form of this particular output, which goes to the GCL via the WSDS, might be as follows for a prototype Red Agent:

Command	**Plan**
Global Command Level	RGCL3
European Command Level	REUR3
Northwest TVD	NWTVD1
Southwest TVD	SWTVD1
Western TVD	WTVD3
Space	RSPA1
Southeast TVD	SETVD2
Intercontinental	IC1
[Others]	

That is, the output is a *vector* of plan names. This might appear in code as a series of Rand-Abel statements such as "Let Plan of Global-Command-Level be RGCL3."

2. *Control variables.* The NCL must also specify controls on the plans chosen. This is essential because the plans are *parameterized building blocks* constructed without knowledge of precisely when they will be used in a simulation. For a moderate set of analytic war plans to be applicable to a broad range of circumstances they must obviously be parameterized in many dimensions. In particular, as we noted above, the NCL is awakened by subordinates on the basis of wakeup rules contained within the subordinates' war plans. Different Ivans and Sams should have somewhat different wakeup rules and those rules should also depend on circumstances. Thus, as a minimum, we require that the NCL have logic to set parameters in generic wakeup rules at the time it specifies war plans. We shall give detailed examples of this in a later section dealing with the prototype models. Setting the control variables is straightforward, with Rand-Abel statements such as "Let control-variable be Yes."

The NCL should have other controls over the plans as well. These might involve authorizations (e.g., the authorization for a subordinate commander to respond in kind to chemical or nuclear use) or specific instructions such as to exclude the bombing of targets that would result in large civilian casualties.[15]

3. *Perception variables.* The NCL model must specify what models or model assumptions are to be used in conducting look-aheads or in evaluating rules involving models of the opponent, third countries, and combat. It can accomplish this straightforwardly with Rand-Abel statements such as "Let Red's-Blue be RB2." Or, in the case of assumptions about third-country behavior or the laws of war, it can have outputs such as "Let Red's Temperament of France be reliable-ally."

4. *Special interfaces for flag wars.* One might expect there to be no direct interface between NCL models (even ultimatums should go through the GCL's diplomatic channels). However, we have had to allow some exceptions as an expedient.

To explain the problem, suppose that an NCL rule-writer wants to initiate space warfare at some point by attacking some of the opponent's reconnaissance satellites. And suppose further that the Force Agent does not currently have any models simulating space warfare. What do we do? The easy answer is that we could forbid the action. However, that would greatly reduce the value of the NCL modeling, because in practice real-world decisionmakers are highly sensitive to many qualitative issues (essentially symbolic issues) that are hard to model well (Davis and Stan, 1984). These include issues such as whether sabotage is occurring in the homeland and whether there are opponent-induced civil uprisings.

All of this suggests that we use an approach analogous to what human control teams use: an approach that announces certain events (*posts certain* "flags") without attempting to simulate them (this philosophy is related to the notion of *scripted models* discussed originally in Davis and Williams, 1982). With this notion in mind we have developed a preferred procedure and an expedient. **The preferred procedure for handling unsimulated events is as follows:**

1. Imbed the actions at issue *in analytic war plans*, however simple (e.g., a space war plan calling for low-altitude antisatellite

[15]The NCL can also receive information and recommendations from subordinates. This information is communicated as a message. The message is accessed via a Rand-Abel variable in the WSDS (e.g., Recommendation-from-SACEUR could be authorize-nuclear-use). If the NCL rules are sensitive to subordinates' recommendations, then that information is available.

AD-A170 952 A NEW METHODOLOGY FOR MODELING NATIONAL COMMAND LEVEL DECISIONMAKING IN W..(U) RAND CORP SANTA MONICA CA P K DAVIS ET AL. JUL 86 RAND/R-3290-NA MDA903-85-C-0030 2/2

UNCLASSIFIED F/G 5/10 NL

END
DATE FILMED
10-86

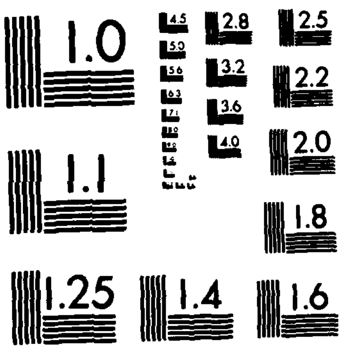

:CROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

attacks). That is, have the flag set by the appropriate agent sending Force Agent an explicit order, just as it would for orders of a sort that would result in real simulation of events.

2. Have the Force Agent post the flag at the appropriate time (e.g, in the event of antisatellite attacks, even a very simple space plan could specify that the attacks would begin at the time of the Soviet invasion of Western Europe rather than at the time the NCL chooses the plan to begin preparations for the invasion, thus preventing Blue from seeing the flag at an unrealistic time).

We have adopted this approach except for the instance of "flag wars" between NCLs in a truncated system without any Force Agent. For that special case, however, the specification must allow for some direct NCL to NCL communication, which otherwise would not be permitted.

LOOK-AHEADS AND PERCEPTIONS

Although the system specifications related to look-aheads and perception models transcend the problem of building NCL models, we discuss them here because it has been in the NCL modeling effort that we have made some of the decisions about how these matters are to be handled. The issue is how to define perception models and how to avoid certain technical problems that could arise from unrestrained use of such models in rule-writing.

As mentioned above, the look-ahead mechanism is one of the novel and important features of the overall RSAC simulation. By allowing decisions to be based on the results of look-aheads, we can significantly increase the intelligence of our Red and Blue Agents.[16] An analogy may be useful: In the real world, decisionmakers reach tentative conclusions based on heuristics and then ask for staffing, which may (or may not) reveal problems or opportunities causing reconsideration. In the same way, our NCL models can use heuristics for tentative decisions and then test them by looking at results of a look-ahead simulation before proceeding.

In principle, one could devise sophisticated heuristics that would do projections without any such game within a game (e.g., some NCL decision rules may ask the Force Agent to run algorithmic models as

[16]The Green Agent could also conduct look-aheads if its rule-writers insisted. However, the current Green Agent rules depend solely on current information. This is not as drastic an assumption as one might think because current information includes, for example, the current rate of advance of armies, the current arrival rate of reserves, and the current sustainability level of conventional munitions—all items providing substantial insight about what tomorrow will bring.

projections). In practice, however, just as in war games played by human teams and real-world decisionmaking, the tendency in developing heuristic rules is to overlook some of the complications and interrelationships that are *usually* not a problem but that can be extremely important in some instances. For example, the current force ratio might be favorable, but a look-ahead might reveal that it would become highly adverse in another week. Although a heuristic might include the prewar planning factor for when force ratios would become adverse, events in the actual war (or simulation) might render that planning factor obsolete.

In any case, look-aheads are an important feature of the RSAC's system architecture.[17] However, they raise some difficult issues.

Obviously, one does not want the NCL conducting the look-ahead to have perfect knowledge of the future. Instead, the look-ahead should reflect *perceptions* of the opponent, third countries, and the laws of war. Although we handle the perceptions of third countries and laws of war by setting parameters (i.e., Red's look-ahead uses Green Agent and Force Agent models with certain parameters changed to reflect Red's perceptions—parameters that may determine which algorithm is used for calculating rates of advance or some such, but parameters nonetheless), we require that perceptions of the opponent be handled *with a distinct* model of the opponent. Thus, we have Red's Blue and Blue's Red as phantom objects in our simulation. Indeed, there may be many possible Red's Blues, and Red must specify which one should be used in a particular look-ahead.

At a qualitative level, this concept seems simple, but attempting to sharpen up the concept reveals several problems:

- If Red conducts a look-ahead using Red's Blue to determine Blue's probable actions, which Red does he use in that same look-ahead? That is, should Red's Red be Red (Red might not understand his own behavior, suggesting that Red's Red is not necessarily Red).
- In the same situation, suppose that in the look-ahead Red's Blue wants to conduct a look-ahead. Does he then use Red's Blue's Red? In instances where Red is attempting a deception operation, for example, would he not expect Blue to have an incorrect image of Red? How far does this recursion go?[18]

[17]Technically, the mechanism of conducting look-aheads involves saving the entire WSDS (pushing onto a stack), running the look-ahead, recording in non-WSDS storage whatever information the rules demand of the projected situation (e.g., the ratio of forces projected ten days hence), and returning to the real game (popping).

[18]The importance of such recursive concepts appears elsewhere. For example, Lefebvre (1982) gives a mathematical-psychology description of differences between ethi-

The potential exists here for the type of infinite recursion that mathematicians are so fond of. In effect, one could imagine starting a simulation and having a look-ahead, a look-ahead within a look-ahead, etc., with no results ever coming back! Furthermore, even if the look-ahead terminated after some depth, there would be an enormous proliferation of rules to be written (those for Red, Red's Blue, Red's Blue's Red, etc.).

Not surprisingly, it is possible to avoid these difficulties with reasonable assumptions. In particular, we note that the marginal quality of information gained goes down rapidly as the depth of such recursions increases—after all, we are not dealing with the highly formalized and rigid rules of chess, but rather with war games and simulations in which one's actions should depend to some extent on one's *estimate* of opponent strategy. There are significant uncertainties about that, but when discussing one's estimate of the opponent's estimate of one's strategy, the information content is modest and can reasonably be stated with heuristics.

Although there is nothing in the software to prevent deeper recursions, we have addressed this class of problems by decreeing that Red's Blues and Blue's Reds do not use look-aheads to make their decisions. That is, their decisions must be made solely on the basis of heuristic rules and calculations not requiring a game within a game. We similarly decreed that Red's Red is Red and that Blue's Blue is Blue, although again this is more of a restriction in procedure than in system software, since it would be straightforward to generate additional subprograms to represent Red's Red 1, Red's Red 2, and so forth. Our reason for *not* doing so is that, once again, we believe the additional complexity would be unjustified.[19]

The construction of Red's Blues and Blue's Reds also poses difficult technical problems. Consider the following:

- On the one hand, it would be reasonable for Red to assume that Blue had observed correctly *most* of his force actions to date—i.e., that Blue's decisions would reflect a generally realistic image of Red's forces, positions, and past maneuvers.
- On the other hand, allowing Red's Blue to have perfect knowledge of Red's actions to date would seem unrealistic given

cal cultures. That description depends on the interactions between pairs of people, including interactions dependent on perceptions to the depth equivalent to our Red's Blue's Red.

[19]It is easy to treat exceptional cases by writing assumptions directly into the rules. That is, Red could do a look-ahead specifying not only that his opponent but also that he (Red) would respond to certain Blue actions in a particular way. What we have tentatively ruled out is formalizing that process by creating Red's Reds.

the so-called fog of war, the interest of both sides in deception, and the possibility for intelligence failures more generally.

- These issues are not merely academic because to conduct a look-ahead simulation it is necessary to specify what world state each side sees.

The baseline. The real question here is what the *baseline* should be: complete ignorance, or complete knowledge. Upon reflection, we concluded that the best approach to take, at least for our *initial* work, was as follows:

- Red's Blue and Blue's Red have complete knowledge of events up to the present time except on specific issues treated explicitly by the analyst by exception (and there may be many exceptions to deal with specific issues of deception or intelligence). Thus, but for the exceptions, Red's Blue knows that Red has been operating on a particular analytic war plan and vice versa. To put it differently, Red's Blue has behaved precisely like the real Blue up until the time the look-ahead begins.
- However, Red's Blue and Blue's Red may show very different behavior than the real Blue and Red, respectively, when looking forward in time—simply by virtue of different NCL rules.

To summarize, then, except when the analyst provides specific assumptions to the contrary, initial **Red's Blue and Blue's Red differ from the real Blue and Red, respectively, only with respect to their NCL components.** Typically, the Red's Blue NCL is relatively simplistic—just as in the real world one tends to reduce one's model of one's adversary to relatively stark terms in most instances.

It would be possible to create specific war plans for use only in look-aheads (e.g., Red's image of Blue's plan for a conventional forward defense of central Europe). However, for the time being, uncertainty about what plan an opponent may adopt provides us an adequate range of possible opponent behaviors. We have not found it worthwhile, relative to our other concerns, to pursue this possibility but it will become important in future work.

Some exceptional cases. We are, of course, interested in perturbing these assumptions to reflect deception operations. For example, Red might believe that Blue would fail to see his covert dispersal of SSBNs and strategic aircraft in a 12-hour period of extremely bad weather. There is nothing to prevent specifying such assumptions as part of the look-ahead procedure. The point, rather, is that one would have to make any such assumptions explicitly as exceptions. Thus, before starting a look-ahead, the Red NCL program would have to have code ordering a change in the Force Agent's disposition of forces relative to "truth." Those orders would be automatically rescinded at the end of the look-ahead. In some instances, such orders would be easy to write. In other cases there would be serious difficulties because of the simulation details in Force Agent—details that NCL plan writers might not know.

Another example of an exceptional case would be our treatment of a Soviet launch of an intercontinental nuclear strike under conditions where U.S. early warning systems had been neutralized. In this instance, if Blue conducts a look-ahead using Blue's Red, he should *not* see Red's launch (i.e., Blue's Red should not launch) until after weapons have landed on the United States, after which Blue would surely "know." The reader will appreciate that attempting to model such asymmetries of information, and updating the degree of ignorance, would be an extremely difficult task if attempted in general rather than for specific issues.

V. THE PROTOTYPE NCL COMPUTER PROGRAM

BACKGROUND

Overview of the Approach

The preceding sections described system specifications in rigorous but relatively general terms. This section describes a particular program *implementation*, relying primarily upon substantive arguments as rationale and upon illustrative Rand-Abel code for examples. Our purpose here is not to document or rationalize a particular NCL model (e.g., Ivan K), but rather to document the program superstructure and rule-writing framework and to convey to the reader a sense of what these models are like and how easily they can be modified.

The principal issues in program *implementation* are, from the viewpoint of computer science: (a) identifying the particular objects to be distinguished (this corresponds to deciding what theaters and command structures to work with); (b) organizing strategic knowledge (i.e., developing useful levels of conflict and useful hierarchies of variables); (c) establishing a decision style (including the style of "search") to be reflected in program logic; and (d) establishing methods for permitting the NCL models to *learn* in the course of the simulation.[1]

Although any of the choices we made in this regard could be changed rather easily within the general system superstructure discussed in the last section, the reasoning/knowledge framework we have adopted is rather robust, allowing us to build alternative Ivans and Sams with a broad range of temperaments and grand strategies. We shall investigate alternative frameworks to some extent, but we believe the greatest intellectual gains can be made by holding the prototype's framework reasonably constant while working very hard to develop substantially richer decision rules.

The prototype program, then, is based on the process model and hierarchy of variables described in Sec. II (see Figs. 2.4 and 2.8–2.9), and the state-space framework described in Davis and Stan (1984) (see also Fig. 3.2).

[1]Ordinarily, choosing a programming language and computing environment is a major decision in implementing a model, but in our case that decision had already been made for the overall RSAC simulation. We use the Rand-Abel language described in Sec. IV (Shapiro, Hall, Anderson, and LaCasse, 1985a, 1985b).

Using the process model has several advantages for NCL modeling above and beyond the substantive basis it has in decisionmaking theory and cognitive sciences. In particular, it provides a natural structure for "chunking" or modularizing the knowledge composing the model and thereby improving the code's overall comprehensibility substantially. Moreover, because the process model's framework is robust, we can create a new model by starting with a skeleton and "filling in the blanks" with rules (knowledge) characteristic of the new Ivan or Sam.

An Aside: Relationships with Other Artificial Intelligence Applications

This subsection is an aside for readers trying to relate our work to the literature of artificial intelligence (AI).[2] Others may skip it entirely. One reason for including this subsection is that **our work is unique or nearly so in the degree to which it: (1) applies and extends AI techniques to a large and complex policy problem, (2) combines rule-based modeling with traditional simulation, (3) relates AI models of different types of behavior (what we have termed "strategic" and "organizational"), and (4) deals with the problem of adaptive simulation.**[3]

Throughout this work we have drawn heavily, and extended considerably, paradigms and techniques from artificial intelligence research generally and work with so-called expert systems specifically (Davis, 1984). For example, our analytic war plans are a hierarchical version of *scripts* and our organization of knowledge in the NCL models is similar in some respects to that of the EMYCIN program. Moreover, the general character of our decision rules emphasizes heuristic symbolic reasoning rather than quantitative algorithmic reasoning (although there are examples of that as well). We have been heavily influenced by the writings of Herbert Simon (see, for collections of papers, Simon, 1980; and Simon, 1982).

As so often happens in practical applications of artificial intelligence ideas, however, we have managed, by focusing on the substantive problem, to avoid many of the classic unsolved difficulties that permeate the academic literature and much of the AI jargon. Because the NCL

---

[2]Two good sources to the literature with extensive bibliographies are Hayes-Roth, Waterman, and Lenat (1983), and Charniak and McDermott (1985). The latter is an AI textbook, which treats expert systems only briefly (pp. 437, ff). The former focuses on expert systems. The research described in the present report involves techniques from expert systems, concepts from AI more generally, a great deal of non-AI computer science, and substantive work in the realm of strategic analysis.

[3]For a rare discussion of the challenge in simulating large-scale complex systems, see Zeigler (1984b). See also Davis (1985b).

models must solve a specific class of problems, we are freed from the rigors of attempting to implement a general decisionmaking system. The strategic thinking in our prototype model is highly structured, and the process model along with the concept of establishing preference orders for analytic war plans allows us to greatly simplify the problem of "search."

Lapsing temporarily into the jargon peculiar to artificial intelligence research, our NCL programs use a high degree of *meta knowledge* in their control structure. This meta knowledge (e.g., the process model and the hierarchy of variables) allows the resulting programs to be implemented *procedurally*. Thus, we do not require a general-purpose artificial intelligence language with the capability for inference and automatic searching.[4] Instead, we specify in some detail the problem-solving strategy, which depends heavily upon heuristic reasoning, and are able to use an extremely fast English-like programming language (Rand-Abel). Although we have a separable *knowledge base*, it contains only the knowledge most fruitful to vary in analysis. Our knowledge base is not literally separated as "data," as in some systems, but rather is accessible to the user by virtue of the ease with which Rand-Abel can be read and changed.

TOP-LEVEL PROGRAM DESCRIPTION

With this background, let us now describe the prototype program, for which Fig. 5.1 gives the top-level structure. The names in the various boxes are indeed the names of particular program modules (Rand-Abel functions). Also, this structure incorporates the NCL process model and the hierarchical situation assessment discussed in Sec. III.[5]

Having provided this top-level view, let us now go through the program from beginning to end, discussing the principal assumptions and providing examples of the decision rules used in the prototype models. Some of the discussion will overlap that of previous sections, but we include it for completeness. The redundancy is a penalty for separating discussion of requirements and specifications from discussion of the implementing computer program.

[4]Most academic work on artificial intelligence is accomplished with LISP (in the United States), PROLOG (in Europe and Japan), or derivatives thereof such as Rand's powerful and English-like ROSIE language (Fain, Hayes-Roth, Sowizral, and Waterman, 1982). These languages have many special capabilities for inference and search but are also quite slow—except, perhaps, for the newer LISP dialects operating on LISP-specialized machines.

[5]The prototype NCL programs are each approximately 4000 lines in length (in Rand-Abel) and are of interest primarily to Rand analysts preparing to develop richer programs for applications work.

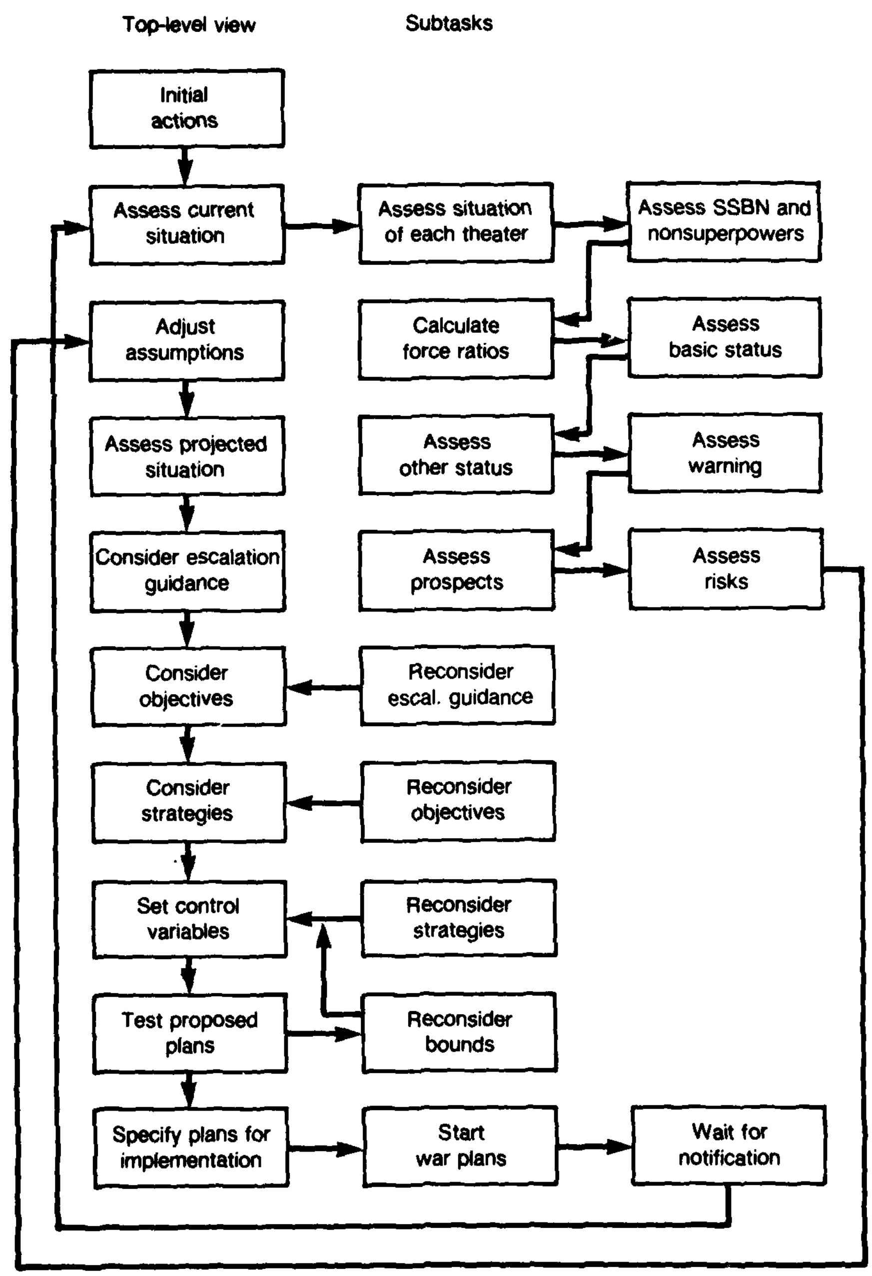

Fig. 5.1—Top-level program structure

INITIALIZATION

Starting up a new game is a technical process of no particular interest to the analyst. The NCL program has certain functions to perform here because of the details of the way in which coprocesses are managed in the RSAC's overall system. In particular, the NCL must install a wakeup routine for itself (which is the only means by which it will ever wake again) and create coprocesses for all of its subordinate command levels. The latter function is accomplished hierarchically by having the NCL create the GCL coprocess, which in turn creates its immediate subordinates, and so on.

Which analytic war plans are used initially is determined by variables in the World Situation Data Set. These would be peacetime plans in most instances, but the analyst can intervene using the Data Editor to invoke other initial plans instead. Once the initialization has been accomplished, the NCL sleeps and awaits some event to awaken it.

WAKEUPS

As with any other coprocess in the RSAC, the NCL will execute only when certain conditions checked in its wakeup routine occur (the tests themselves are performed by the Force Agent and System Monitor on a regular basis). The NCL prototype models may awaken for any of the following reasons:

1. Because of military or political events in the simulation, there is need for an NCL decision, which the lower command levels request.
2. One NCL model (say, Ivan K) is conducting a look-ahead test, which requires another NCL model (e.g., Ivan K's current model of his opponent, that is, Red's Blue 1) to awaken and make a decision.
3. A human analyst (or team) wishes to induce a plan substitution.[6]

Of these, the first may be considered the standard case and the others "special." In the special case of a look-ahead by, say, Red, the NCL model at issue is Red's Blue NCL rather than the real Blue NCL.

[6]There is actually another case, one beyond the scope of this report, in which a human team wishes to run the NCL model as a source of advice without having the NCL's recommended decision being directly implemented. This decision-aid function is under development currently.

Both reside in the same NCL coroutine. As mentioned in Sec. IV, when Red's Blue NCL program executes (or Blue's Red NCL), there is no full-system look-ahead, although there may be rules reflecting some notion of Red's Blue's Red NCL.

In the special case where a human player requests a plan substitution, the substitution is made directly, after which the NCL sleeps without further processing.

SITUATION ASSESSMENT

One of the most interesting features of the NCL models is the way they conduct situation assessments. As discussed in Sec. II, they begin by looking at low-level details (weapon counts, political cooperations, etc.) and then fold such details together to reach judgments about increasingly higher-level issues (see Figs. 2.6 and 2.7). *We emphasize that this situation assessment is by no means "objective," but is rather the result of objective factors, perceptions, and the temperament of the particular Ivan or Sam.* Thus, situation assessment may be affected by, for example, erroneous intelligence information, loss of communications, or the NCL's "mindset."

It may be useful to give a few examples of situation-assessment rules as they appear in the Rand-Abel language within the Ivan K prototype (see the following page).

These examples should convey the flavor of Rand-Abel code and also something of the technique for writing NCL rules. Consider, for a moment, the first rule, that regarding Risks. The rule makes no sense, of course, unless one has previously defined a number of variables, some of them qualitative and some of them quantitative. Warning-of-escalation is defined with rules very much like this one, but in terms of whether Blue's nuclear forces are dispersed and poised for action, whether there is political intelligence about Blue planning a nuclear strike, and so on (matters followed by the simulation).

The variable report-from-price-of-going-second-IC is different in kind. This describes the price (measured as a percentage change in the ratio of intercontinental weapons) of going second in a nuclear exchange rather than first. It is an *example* of information that Ivan might be concerned about as the nuclear threshold approaches. It should be evaluated by a subroutine taking into account the nature of the missile arsenals, accuracies, reliabilities, and vulnerabilities.

In evaluating Risks with the above rule, Ivan K had to combine different types of information (objective indicators of warning and results of a theoretical calculation based on his perceptions of missile characteristics and other factors). The second example, the evaluation of

Assessing Risks

```
If Current-situation  = Eur-gen-tac-nuc
Then
  :

    If Warning-of-escalation  Eur-nuc and
    (the report from price-of-going-second-IC   50.0 [% in COF] or
    Disarming-capability is high)

    Then Let Risks be high.
  :
```

Assessing Main-theater Status

[SCHEMATIC COMMENT

Value of Main-theater-status
as function of Ground-status and Eur-nuc-COF

Ground-status

GM	SL	L	NP	P	GM
P	SL	L	NP	P	P
NP	SL	L	NP	NP	NP
L	SL	L	L	L	L
SL	SL	SL	SL	SL	SL
	SL	L	NP	P	GM

Eur-nuc-COF

]

```
If Current-situation  = Eur-gen-tac-nuc
Then
  :

  If        Ground-status is goals-met
    and     Eur-nuc-COF  = 2.0
  Then Let Main-theater-status be goals-met.
  Else If
            Ground-status  = progress
    and     Eur-nuc-COF  = 1.5
  Then Let Main-theater-status be progress.
  Else If
            Ground-status is serious losses
    or      Eur-nuc-COF  = .65
  Then Let Main-theater-status be serious-losses.
  Else If
            Ground status  no-progress
    or      Eur-nuc-COF  1.0
  Then Let Main-theater-status be losses.
  Else Let  Main-theater-status be no-progress.
  :
```

Main-theater status, shows even better how Ivan K must make tradeoffs between apples and oranges. As the schematic in comments at the top of the module indicates, the reasoning here is that Ivan K will consider his Main-theater-status to be goals-met only if results have been good in terms both of ground gains and maintaining a good ratio of nuclear forces (a good "correlation of forces"). The entries in the schematic show the value of Main-theater-status with abbreviations such as GM for goals-met. The actual code follows, here in the If—Then—Else form, although an equivalent decision table might be used if Ivan K were to be reprogrammed today.

To apply the rules establishing Main-theater-status, the program must already have values for ground-status, which is itself determined by rules (taking into account how far Red's forces have penetrated along various major corridors) and for Eur-nuc-COF, which is determined by an algorithmic definition plus current data from the combat simulation.[7]

It may be interesting to mention in passing that in developing the various decision rules the analyst cannot rely solely upon the criteria to be found in Red and Blue doctrine, since doctrine may represent nothing more than a prewar theory of "norms." For example, it is well known that Soviet doctrine calls for fast advances—tens of kilometers per day during the initial phase of invasion in Europe. However, if the Soviets were to find themselves in war, they would by no means escalate (or terminate) *merely* because their narrow doctrinal performance goals were not being met (although the authors have heard claims to that effect from time to time). Instead, the Soviet commander would have to assess the pros and cons based on current realities. The analyst, then, must draw upon doctrine for insight but cannot depend upon it.[8]

It is probably useful to give one more example of the issues arising in Situation Assessment. Consider now the issue of evaluating "Other-status", a variable characterizing the status of events elsewhere than in the main theater (see Sec. II). Upon reflection, one realizes that a decisionmaker would be interested less in precisely what is happening "elsewhere" than in whether what is happening is worse than

---

[7]More realistic rules for establishing Main-theater-status would probably focus less on correlations of force than on estimated Red and Blue capabilities to accomplish certain military objectives with nuclear strikes if those strikes should be necessary. The simplification here is for methodological convenience in a prototype.

[8]One reason we are loathe to refer to our work as an example of expert systems is that there are no experts in this realm. Unfortunately, one needs a great deal of insight to build useful NCL models, and as indicated in Davis and Stan (1984), there are major discrepancies between what the nations' strategists and doctrinal writers seem to be saying and what they would say if the context of discussion were better defined.

be expected given the main-theater situation. Thus, if the Soviets invaded Southwest Asia and found themselves fighting U.S. CENTCOM and supporting forces there, it is most unlikely that they would be surprised if U.S. forces sank Soviet navy vessels in the Indian Ocean. That would probably *not* constitute a "bad" Other-status. On the other hand, a worldwide naval campaign (as distinct from some isolated maritime incidents) might be viewed with considerably more concern.

With considerations such as this in mind, rules evaluating Other-status might look something like the following (again drawing upon the prototype version of Ivan K):

If the Current-situation is SWA-gen-conv
Then

Decision Table

Current-situation-other-land	Current-situation-other-naval	Current-Situation-USSR-homeland	Current-situation-space	Current-SSBN-warfare /	Other-status
--	Demo-tac-nuc	--	--	--	bad
--	Gen-tac-nuc	--	--	--	bad
--	--	--	Strat	--	bad
No-war	--	No-war	Tac	--	marginal
No-war	No-war	No-war	--	--	good
No-war	Demo-conv	No-war	--	No	marginal
No-war	Gen-conv	No-war	--	No	marginal
--	--	--	--	--	bad.

In reading this table, which uses the condensation possible with use of the "--" feature, one should recognize that the third and subsequent rows never apply unless the current "other-naval" situation is "less than" nuclear because the computer leaves the table once it sees a row for which all the If conditions are true. The fourth row shows the case alluded to above in which Ivan K characterizes the situation as marginal even though there could be global naval war (and some antisatellite actions). We should emphasize yet again that these rules are purely illustrative and apply only to a particular model of the Soviet Union developed for prototype purposes. The last line, which the computer reads only if none of the preceding lines' If clauses are true, is a catchall for other cases. Other prototype models make very different situation assessments. On the other hand, the types of variable emphasized are those identified as important in Davis and Stan (1984).

MODELS OF THE OPPONENT, THIRD COUNTRIES, AND THE LAWS OF WAR

Models of the Opponent

After performing the situation assessment, the prototype NCL programs update their assumptions about opponent, scenario, and force—assumptions that are reflected both in ordinary rules and in rules requiring look-ahead projections. The technical issues involved in developing opponent models and corresponding programs were discussed in Sec. IV and will not be repeated here.

Models of Third Countries and the Laws of War

By analogy with our need for a Red's Blue, we include the concepts of Red's Force and Red's Green different from the true Force and Green Agents. That is, Red will make predictions about events in the war and decisions by third countries based on his own models of combat and third countries—models which may or may not agree with what the simulation regards to be "truth."

In implementing Red's Force and Red's Green we have *not* developed them as separate programs. Instead, because the Force Agent and Green Agent are highly parameterized, it is possible to invoke Red's Force by merely having Red specify some parameter values to be used in running the Force Agent. The same process applies for Red's Green. Thus, at the beginning of the game, one must specify not only the temperament of Saudi Arabia but also Red's and Blue's perceived temperaments for Saudi Arabia. If the Red Agent wants to make a projection using a different assumption, it is necessary only to write the corresponding rules. In principle, analysts developing Red Agents could employ look-aheads varying any or all of the hundreds of parameters contained in the Green and Force Agents. In practice, we assume nearly all of the parameters are the same for Red, Blue, and "truth's" versions of Force and Green—focusing on only a few key variables of interest to the particular analysis.

Adjusting Assumptions: Learning

On each waking, the NCL has the opportunity to modify its assumptions based on the events that have taken place in the game. This updating represents a limited but important form of *learning*. There exists no general theory of learning for use in artificial intelligence models, and it is unlikely that computers will display the full range of human-like learning for a very long time, if ever.[9] Nonetheless, we have included certain important learning features in the NCL model framework and have written illustrative learning modules for the current prototypes. The basic notion behind our approach is that it should not be particularly difficult to have the models learn in the course of a simulation if what they are asked to learn is something they are prepared to learn, or something they "almost" understand in the first place.

Probably the best approach here is to give some concrete examples of possible learning logic.

1. *Opponent model.* As we have discussed previously, an important feature of the overall simulation and of the NCL decisionmaking in particular is the assessment of one's opponent. Indeed, that has been a dominant factor in some important historical decisions such as Hitler's decisions about campaigns in Poland and France. Even if one has a strong initial image of the opponent, however, the opponent's *actions* once mobilization or combat begin must obviously be taken into account. This is especially so if the possibility exists that the opponent's government may be changing.

How difficult it is to build logic allowing an NCL model to learn about its opponent—or to change its opponent model altogether—depends on how sophisticated Red's Blues and Blue's Reds are. Even very simple versions of such constructs add substantially to the value of game-structured simulations. To illustrate, consider six Red's Blues with the following top-level characteristics:

Descriptions of Alternative Red's Blues

- Blue 1 will not escalate further and will tend to be slow in reacting and indecisive generally, although he would eventually mobilize and fight for regions he considered vital. Under some circumstances, Blue 1 would consider backing off, disengaging, or even surrendering rather than risk a large war.
- Blue 2 will try to follow traditional incrementalist Western thinking: He will escalate up to and including demonstrative use of nuclear weapons in Europe, if necessary and feasible, but will not go beyond that. He will be faster to act and more nearly decisive than Blue 1 but will still seek a fair degree of consensus before acting. He may escalate war to the high seas. (Blue 2 is close to the Sam 5 we have used in our illustrations.)

[9]See Schank (1984) for a highly readable popular discussion by one of the pioneer researchers in AI. The book is notable in part for its candor and lack of AI hype.

- Blue 3 is similar to Blue 2 except that he will follow the declaratory policy of NATO planning to the limit—launching a general single integrated operations plan (SIOP) in coordination with the NATO strike plan when and if the time comes. Blue 3 is therefore more resolute but is not unilateralist. If he launches the SIOP it will be a full countervalue attack.
- Blue 4 is substantially more unilateralist and decisive; he will follow the spirit of NATO planning but will generally be faster to act (not requiring consensus). He will launch the SIOP in connection with a NATO strike plan. He may execute major portions of the plan without full NATO approval. He will not try to create a homeland sanctuary. In Ivan's view, Blue 4 is irrational.
- Blue 5 is much more unilateralist and decisive but is also more pragmatic. He will not endanger the U.S. homeland if that can be avoided. Thus, he will be faster and more decisive at levels of conflict up to and including European counterforce strategic warfare in Europe but will not launch the SIOP (regardless of NATO planning) unless he deems it essential for U.S. survival.
- Blue 6 is a worst-case Blue in many respects. He is unilateralist and decisive, but is also willing to fight nuclear wars to avoid losing. Thus, Blue 6 is capable of both preemption and first strikes at almost any level.

In thinking about these Red's Blues the reader should remember that it is irrelevant that real national governments are more complex in their decisionmaking than these short descriptions suggest. After all, the Red's Blues are merely *planning constructs* for Red, and there is a good deal of evidence from psychological studies that our cognitive structures of objects are greatly simplified versions of what they represent.

With this background, then, let us consider briefly what might be entailed in allowing Ivan K to "learn" from events when deciding what assumptions to make about his opponent. Suppose Ivan K invades Southwest Asia believing the United States will not even show up for battle (i.e., his initial model is Red's Blue 1). Suppose, however, that after a sequence of events conventional war begins in Europe. How does Ivan K reassess his image of Blue? The logic might be as straightforward as what follows:

Decision Table

[--]

Current-situation	Warning-of-escalation	Time-since-D(Eur)	Presumed-opponent /	/ Presumed-opponent
Eur-gen-conv	None	long	--	Blue1
Eur-gen-conv	None	short	--	Presumed-opponent
Eur-gen-conv	Eur-nuc	--	Blue 1	Blue3
Eur-gen-conv	Eur-nuc	--	>Blue 1	Presumed-opponent
Eur-gen-conv	>Eur-nuc	--	--	Blue6.

[long means greater than 10 days]

[The last rule reflects the nonpragmatic aspect of the case; hence, Blue is not Blue 5 and Ivan K now assumes the worst-case Blue 6.]

The reader will quickly appreciate that far more complex rules could be constructed.[10] However, note also that much of the work in making this decision about probable opponent is done in the evaluation of warning—which we do not discuss here. How early Ivan regards the various indicators as indicating (serious) warning of a Blue escalation depends on Ivan's temperament and his judgments about what actions by Blue are prudent preparations and what actions are truly provocative.

2. *Third-country models.* Similar techniques can be used to adjust NCL assumptions about third-country behavior. This is absolutely essential, not merely nice to have, because we want to conduct games and simulations varying the participation of such nations as Poland, Saudi Arabia, Belgium, and France. If the Ivans and Sams start with assumptions about the reliability of their allies, they clearly *must* adjust those assumptions if their allies decide to sit out the war. If they did not make these adjustments, they would—in the straightforward approach taken by computers—continue to issue orders, proceed to war, and count on reinforcements as though their allies were participating.[11]

[10] Note also some technical elegancies of this table, including the use of the variable "Presumed-opponent" in three ways: (1) as an independent variable, (2) as a dependent variable, and (3) as an indirect value of a variable meaning "do not change this value."

[11] This is not entirely true in the RSAC simulation, since the military command levels of the Red and Blue Agents make most decisions about allocation of forces on the basis of the forces available to be allocated. Thus, they will fill vacuums without the NCL telling them to do so. They may even notify the NCL that he should reconsider his decision about an analytic war plan because allied forces are not showing up. However, if the NCL model is not smart enough to pay attention and revise his assumptions, he may well stubbornly insist that the plan be pursued without modification.

Actually writing the decision rules for this learning process can be more tedious than writing rules for the opponent models, only because there are so many third countries. However, by providing "packages" of assumptions at the right granularity, tables like those for assumed opponent can be developed (although we have not yet done so).

3. *Laws of combat.* All too frequently, strategic analysis proceeds as though our models of combat are correct—even though we know, with certainty, that they are likely to be seriously wrong in at least some respects. So, for example, there are vast uncertainties regarding the likelihood of early Soviet breakthroughs in Europe (even with all force levels specified), rates of advance with and without breakthroughs, and attrition levels. And, similarly, it would not be surprising if entire weapon systems failed in the event of general nuclear war. In the past, we have discovered gross problems of reliability in deployed strategic systems and there can be no assurance that there do not exist similar problems today.

Given the large uncertainties, it is not surprising that Soviet and Western planners make different assumptions in their calculations about how war will probably go. Some of this is at the level of strategic assumptions, and some is at the level of functional forms or parameter values within those functional forms used to project movement rates, kill probabilities, and the like. For example, it is possible that the Soviets project rates of advance in the first days of a European conflict that would exceed by factors of five to ten the rates of advance predicted by best-estimate U.S. models.

Although we have not yet had time to explore mechanisms of learning in this realm, it is easy enough to see places to start. Taking the example of Soviet rates of advance in Western Europe, it would be possible for Ivan to do a calculation based on his empirical movement rate and adjust the key parameters in those calculations accordingly. This would not be entirely straightforward, because there are many parameters in those expressions, thereby requiring "intelligence" in making the adjustments. If Ivan's concept of what might go wrong is correct, his learning will be valid; if what has gone wrong is different in character, his adjustments may make things worse rather than better.[12]

---

[12]To illustrate technically what can happen, suppose that Ivan uses an algorithm for projecting movement rates that depends only on force ratios (not uncommon in modeling). If he observes zero movement in the first few days of war, he might—stupidly—infer that the parameter called Vmax in his algorithm should be set to zero (or to something approaching that, after weighting a priori estimates and current empirical estimates in some Bayesian manner). In fact, what might be happening is that the functional form of his laws of combat is wrong: The correct expression might predict that

This discussion has probably been sufficient to convince the reader of the following:

- Learning is *essential* in the types of simulation we are discussing.
- Some types of learning can be modeled straightforwardly, with the potential for incremental improvements in sophistication. These types of learning require the Ivan or Sam to have a reasonable a priori understanding of mechanisms and alternatives (e.g., a list of plausible Red's Blues or an algorithm for predicting movement rates that is reliable except for some parameter values).
- Other types of learning can be imagined but are beyond what we have attempted so far. These include: (a) using multiple regression techniques to infer laws of war as the war continues; (b) using rule-based logic and early results of crisis or conflict to build up new models of the opponent or third countries using, for example, standard attributes such as those discussed in Sec. III; and (c) providing the models with much richer sets of possibilities to explore before making adjustments (e.g., more Red's Blues and more complex algorithms). All of these would require major efforts.
- More realistically, we believe that much can be done with the types of learning that are within our grasp *and* a pragmatic attitude about man-machine interactions. It is *not* essential, nor even particularly desirable, that the NCL models or the other decisionmaking models of the overall RSAC simulation be entirely self contained. To the contrary, it is reasonable, desirable, and inevitable that more and more features will be added to the models as we gain experience from games and simulations exploring nonstandard cases.

LOOK-AHEAD PROJECTIONS

Having characterized the current situation and adjusted assumptions, it is natural to project ahead to see what the future is likely to bring if one continues on the present course of action. With this in mind, we include in the NCL program *automatic look-aheads*, which provide information about this projection of the future. This

---

the defense should be able to hold with zero movement if ordered to do so and if provided with enough forces to maintain a certain density of forces on line—independent of force ratio.

information can be used in rules without the rule-writer having to define the code for look-aheads. In some instances, the decisionmaking process of the NCL does not use any such information, in which case this look-ahead will not be performed. In some instances of course, the automatic look-ahead's results are specious because they assume continuation of the current plan when such an assumption is silly on its face. Nonetheless, *as a rule*, it is natural to expect decisionmakers to want to know "What will happen if we just press on with current plans?" And in the case where the current plan is clearly not viable, the results of the look-ahead will indeed be negative in any case.

Technically, performing the automatic look-aheads raises several issues. The first of these we have already discussed—the issue of what assumptions to make. Our *automatic* look-aheads use best-estimate models of the opponent, third countries, and combat. A second issue is how to make the look-ahead work—after all, the reason the NCL program is operating is that there has been a wakeup by lower-level commanders, usually because things are going wrong with the current plan. In that case, the NCL is awake because a "bound has broken" (wakeups are triggered by rules that continuously, or at least systematically, test to see if certain conditions invalidating the plan have been met). If the look-ahead literally uses the same plan, it will break the same bound again.[13] The solution, of course, is to turn the bounds off when performing a look-ahead.

Another issue is knowing what to do with the information generated in the look-ahead—what should be stored, what should be processed? Currently, our approach is to examine the state of the world at two points—three and 14 days. The state of the world at those points can be stored and referred to directly in NCL rules. In addition, the rule-writer may insert logic to accomplish situation assessments at the two points; if he does so, then information such as the Main-theater-status projected for three and 14 days in the future will be available as variables.

In practice, we have made little use of this facility in the prototypes, primarily because the first-generation NCL models were developed before the full integrated system was operating properly, and also because it was difficult enough—in a first attempt—to organize information based on current information.[14]

---

[13]If one part of a program calls another, which then calls the first, which then calls the second, and so on, the situation is described picturesquely as a "deadly embrace."

[14]Although look-aheads are a very powerful tool, for both human teams and automated agents, current information plus general knowledge of some relatively simple algorithms can substitute for them quite well in most cases. For example, if the defender is outnumbered 5:1, it is not essential to conduct a full-system look-ahead to know that

MAKING DECISIONS ON ESCALATION GUIDANCE, OPERATIONAL OBJECTIVES, AND OPERATIONAL STRATEGY

The previous modules of the NCL program deal with initialization, situation assessment, and projections—all related to *information*. By contrast, the modules on escalation guidance, operational objectives, and operational strategy deal with decisions of first-order importance, decisions to be made largely on the basis of the higher-level variables evaluated in the situation assessment—i.e., variables such as Status, Prospects, Risk, and Warning.

The decisions may depend on the particular temperaments of the agents, with some tolerating higher risks than others. Also, some will have different value schemes for trading off the significance of events in different theaters. And, fundamentally, some will be more averse to the use of nuclear weapons than others.[15]

In addition, it should be emphasized that the situation assessments accomplished before the decisionmaking steps are not "objective" but rather the result of looking at objective information (some of which could be erroneous) through the lens of a particular temperament. Thus, different Ivans might well evaluate the identical situation as good, marginal, or bad. Once again, the reader is cautioned that the various decision modules are only formally independent. In fact, the rule-writer must maintain coherence across modules, which means that there is *implicit* correlation. So, for example, a decision by Ivan 3 that depends on the value of Risks will depend on the value of Risks having been determined by Ivan 3's rules.

As explained above, the decision on escalation guidance is made with rules reflecting the agent's overall grand objectives and grand strategy, which are part of what characterizes that agent (see Sec. III).

Shown below is a representative decision table for the decision on escalation guidance. In this, the particular Ivan is *not* behaving like the purely doctrinal stereotype (although we tend to believe that U.S. interpretations of Soviet doctrine are often oversimplified). Even though there is demonstrative use of tactical nuclear weapons in

---

prospects are bleak. One problem with this approach, however, is that one tends to overlook things. For example, in the case of a 5:1 disadvantage, one might forget to ask whether reinforcement streams will change that ratio. It is not at all unusual in games and simulations to have such ratios change dramatically (e.g., from 5:1 to 2:1) because of something like "Sealift arrives on Day N."

[15] We plan to experiment with generic NCL models parameterized by the temperament attributes of Sec. III and, probably, simple characterizations of their military grand strategy. The strength of the parameterization approach is also its weakness: It proliferates the number of variables that need to be set.

Europe—a use that some would argue would immediately lead to full-scale nuclear use by the Soviet Union—this particular Ivan first examines the possibility of succeeding with plans that do not further escalate. (They might include tit-for-tat exchanges, but no general use of tactical nuclear weapons in Europe, much less an attack on the U.S. homeland.) On the other hand, if such plans fail look-ahead tests, then he will consider escalatory plans. The mechanism for this is the *preference order method* mentioned above. Note that the decision table specifies not only the escalation guidance, but the second-choice guidance if plans consistent with the first guidance fail in look-ahead tests.

If Current-situation is Eur-demo-tac-nuc
Then

Decision Table

Main-theater-status	Other-status	Prospects	Risks	/	Preference-order	Escalation guidance
goals-met	--	--	--		1	Terminate
progress	good	good	low		1	Eur-demo-tac-nuc
progress	good	good	low		2	Eur-gen-tac-nuc
progress	marginal	good	low		1	Eur-demo-tac-nuc
progress	marginal	good	low		2	Eur-gen-tac-nuc
progress	good	good	marginal		1	Eur-demo-tac-nuc
progress	good	good	marginal		2	Eur-gen-tac-nuc
progress	marginal	good	marginal		1	Eur-demo-tac-nuc
progress	marginal	good	marginal		2	Eur-gen-tac-nuc

[many other cases, truncated for brevity]

End.
}

Generally speaking, we believe it is reasonable for escalation decisions to be decided largely on the basis of high-level variables. Certainly, human teams in war games deceive themselves if they believe they are really using much of the detailed information available in such games when they make escalatory decisions. By contrast, decisions on operational objectives and operational strategy must, of necessity, involve more detailed information. Consider, for example, some possible variations in operational strategy for a Soviet invasion of Western Europe. Variations might involve:

- Premobilization preparations (training of low-readiness units, perhaps over a period of a year or so);
- Nominal mobilization time before invasion;
- Criteria for changing the invasion time in response to Blue preparations for war;
- Alternative use of non-Soviet Warsaw Pact forces, forces of uncertain quality and reliability;
- Alternative concepts of maneuver (massing on 1, 2, 3, or more axes of advance);
- Dependence on operational maneuver groups to achieve early breakthroughs;
- Conducting a massive air operation deep into NATO territory (potential for attrition of NATO nuclear forces but also for substantial losses of air power);
- Operations on the northern and southern flanks of NATO;
- Alternative use of the Soviet navy worldwide; and
- Possible feints in other areas, perhaps by Soviet surrogates or by nations influenced by the Soviets (e.g., Cuba or North Korea).

The list could go for pages. Our point is that the decision rules for choosing one operational strategy rather than another will necessarily involve detailed information such as how far into mobilization the various NATO allies are: the state of NATO's defensive preparations with respect to mining, barriers, and the like; projected results of key battles (e.g, the battle for air superiority or the first-day's battle in NATO's weakest sector); and so on.

Currently, the prototype NCL models do not contain the richness of detail that is possible and essential for applications work. Although improved models are under development for use in studies, it would be premature to elaborate further in this report. Suffice it to say that from the standpoint of methodology, the same techniques apply as in the prototype models.

One point is important to make here, however. In selecting operational objectives and strategies, the NCL must make decisions for all theaters of the world. Thus, as indicated in Sec. IV, the output of an NCL decision in our prototype models consists of a *vector* of plan names as shown below for a representative Red decision:

Illustrative Output (Analytic War Plan) of NCL Decisionmaking

Global Command Level	RGCL3
European Command Level	REUR3
Northwest TVD	NWTVD1
Southwest TVD	SWTVD1
Western TVD	WTVD3
Space	RSPA1
Southeast TVD	SETVD2
Intercontinental	IC1
[...others...]	

The plans, then, are ultimately chosen by name. To each plan (e.g, RCL3, which is a Global Command Level plan for coordinating actions across theaters), there corresponds a complete program embodying the associated escalation guidance, operational objectives, and operational strategy relevant to the command level at issue. These programs also have numerous parameters providing us with considerable flexibility in adapting generic plans to different circumstances. To give an obvious example, the nominal mobilization time before a Red invasion of Western Europe is a parameter that can be changed by the NCL (see below) or the analyst.

SETTING CONTROL VARIABLES

Once the NCL has chosen escalation guidance, operational objectives, and operational strategy, it has chosen an analytic war plan. However, that analytic war plan is typically parameterized in many ways to allow it to be used under diverse circumstances. The parameters include the *bounds* determining when the NCL is to be awakened, *authorizations* for the military command levels to take certain actions on their own, and certain others such as the mobilization time that a given Ivan might insist upon before launching an invasion in Europe.

If the analysts writing NCL models had to concern themselves with all the intricacies of the parameters available to them, it is likely they would be paralyzed rather than enthusiastic. Many of the parameters, after all, deal with detailed operational-level issues. It is therefore essential that the analysts developing analytic war plans have a reasonable understanding of what the NCL models would be expecting (just as it is important in the real world that operational rules of engagment take into account concerns of the national leadership). In our simulations and in the real world there will be disconnects. Indeed, studying those disconnects may prove to be highly valuable.

The mechanism we use for handling such problems is to have the parameter settings exist in the NCL program, but with most NCL rule-writers copying over default values from a standard file after only a casual review to assure consistency with their NCL's philosophy. Thus, our approach here is that the NCL programs *can* specify actions in substantial detail, but ordinarily will not.

To illustrate how we establish default values for authorizations, consider the following table, which establishes authorization for theater commanders to use nuclear weapons if they reach a point in their campaign where their analytic war plan calls for them to do so, given prior permission.

Decision Table

Escalation-guidance	/ RGCL-a	REUR-a	WTVD-a	NWTVD-a	SETVD-a	Space-a	IC-a
Prep	No	No	No	No	No	No	No
[skipping many levels for brevity, with No's appearing in all cases]							
Eur-demo-tac-nuc	Yes	No	Yes	No	No	No	No
[skipping more levels for brevity]							
IC-strat	Yes	No	Yes	No	No	No	Yes
--	No	No	No	No	No	No	No.

Let Nuclear-authorization of RGCL be RGCL-a.
Let Nuclear-authorization of REUR be REUR-a.
[and so on]

This particular set of rules has more of a computer flavor than some, for technical reasons we will not discuss here. Suffice it to say that the default rules in this instance specify that when the level of conflict reaches that of demonstrative nuclear use in Europe, then (in this particular example, which is debatable) the particular Ivan authorizes nuclear use by his global military command and by his commander responsible for the Western TVD—Westerners would call this the European central front. He withholds authorizations elsewhere, however. Also, unless the NCL decision specified escalation, any use authorized by this rule would be strictly limited in nature—e.g., a response in kind.

There are many bounds and authorizations to be considered, only some of which have meaning in the prototype systems. The bounds at issue involve nuclear use against the homeland, strategic force dispersal, mobilization by the opponent, escalation by the opponent, initiation of combat, major changes in political alignments or cooperation, and so on. The authorizations involve alerts, dispersals, mobilizations, nuclear and chemical use, and a number of others.

PLAN CHECKOUT AND THE RECONSIDER FUNCTION

At this point, then, the NCL program has established an analytic war plan and has tuned it somewhat by setting control parameters. The next step is whether to test the plan (if the final decision has not yet been made) or to implement it altogether.

The reader will recall from Sec. II and elsewhere that our NCL model *tentatively* chooses an analytic war plan, but then tests that plan in a look-ahead before implementing it. The look-ahead serves much the same function as a careful staff study *following a tentative command* decision in the real world. Also, because our rules are far less robust than human intelligence, the look-ahead is a mechanism for assuring that the choice of plan (which was driven by heuristics) was not a bad one merely because of something predictable that real-world staff officers would catch at the time. Sometimes the look-ahead test will fail for reasons that are obvious in retrospect, but that were previously unrecognized. For example, a look-ahead might fail because critical forces were already being employed elsewhere, or because—upon reflection—one would expect the opponent to preempt during one's own preparation phase. Logistics, mobility, overflight and basing rights, and a multitude of other problems could prove the limiting factor.

Upon learning that its tentative plan fails a look-ahead test, the NCL must reconsider. At one extreme, we might have the entire decision process begin again, but with the failed plan removed from the list of candidates. Such an approach would, however, be quite at variance with realistic planning and also quite inefficient. In the prototype models we have written NCL decision rules so that when the NCL picks his plan in the first place, he also specifies, in rank order, the subsequent plans to be tested if the first one fails in look-ahead. In this way we can enforce easily a logic such as: First, reconsider the range of operational strategies; if none of those succeed, reconsider the range of operational objectives; if none of those succeed, then reconsider escalation guidance; if none of those succeed, then notify a human operator of the difficulty. Moreover, we can also capture some of a strategist's *common sense* through this rank-ordering scheme. A more mechanical software scheme might search through the space of possible war plans in any of several ways (e.g., always trying the escalatory paths, or always trying the deescalatory paths, or trying some random walk, . . .).

To illustrate this, which must seem more abstract than it needs to be, consider an incrementalist Ivan considering a possible change to his analytic war plan after his initial invasion of Iran has been met by U.S. CENTCOM forces with a tripwire. Such an Ivan might well choose a plan intended to achieve quick and decisive conventional victory in Southwest Asia (an escalation), but if all his look-aheads (staff studies) indicated the likelihood of failure, it is probable that the incrementalist Ivan would then terminate the war with partial gains rather than consider escalating further (e.g., to nuclear use, and/or to war in Europe). Thus, after examining alternative strategies and objectives, that Ivan would probably consider deescalatory rather than even more escalatory options. On the other hand, there would be circumstances when the same Ivan would escalate as necessary to achieve his grand objectives—without seriously considering war termination.

Our point here is that given a set of analytic war plans, with varied levels and scope of war and with varied objectives and strategy at the operational level, it is not too difficult for an analyst composing coherent rules to rank order the plans for testing. In doing so, however, he is bringing to bear an enormous amount of class knowledge about values, tradeoffs, grand strategy, and the like—for his particular Ivan or Sam. To try to reproduce that knowledge with some kind of generalized software to search through the analytic war plans would be difficult indeed. Furthermore, a goal-directed search testing all the plans to determine which was "best" might prove very difficult in practice because translating conflicting vaguely stated goals into rules for action is notoriously difficult for humans, much less computers. Indeed, real-world decisionmakers do not typically evaluate a large class of options. Instead, they filter (using heuristics of the type we try to capture in rules) and then analyze the remaining contenders in depth. They do not readily go backward and reconsider previously discarded options once a course of action has been tentatively chosen but difficulties encountered.

Although we use a rank-ordering approach in our current prototype models, the general system software permits us to use different decision styles. Notable among these are:

- Comparing two (or more) candidate plans with rigid criteria before selecting either; and
- Testing the candidate plan against two (or more) sets of assumptions and criteria for success: e.g., a best-estimate set of assumptions, with ambitious criteria, and a worst-case set of assumptions, with criteria establishing a worst-acceptable failure.

Although nothing in the general approach we have taken precludes such decision styles, we have not yet developed the specialized software to make them easy with a minimum of code writing. We shall probably insert something like the following logical cases in the NCL code, and then develop simple procedures for implementing Cases A, B, C, and D. The prototypes (and, arguably, many decisionmakers) use Case A.

RECAPITULATION

This concludes our walk through the program structure. The principal conclusions the reader should be drawing are essentially these:

- The program structure is highly consistent with the conceptual model and the system specifications emerging from that model.
- The program superstructure is robust enough to cover not only a wide range of Ivans and Sams, but also a wide range of scenario classes (e.g., classes of scenario focused on the Middle East or Far East rather than Southwest Asia, Europe, and the intercontinental theater) and a wide range of concepts on how a given NCLs should do their "reasoning."
- The framework used to organize reasoning and knowledge is also robust, although not so robust as the purely technical aspects of the program. It should be possible to represent a wide range of alternative Soviet and U.S. behavior patterns within the basic framework by merely changing the individual decision rules. Concepts such as the process model of decision and the hierarchy of decision variables (with high-level variables of status, prospects, and risks) are quite versatile.
- The search approach of choosing and rank-ordering plans using heuristic rules (and such special look-aheads as rule-writers wish to specify), and then selecting the first plan to pass a more standardized look-ahead test, is not only convenient but *consistent* with much real-world decisionmaking.

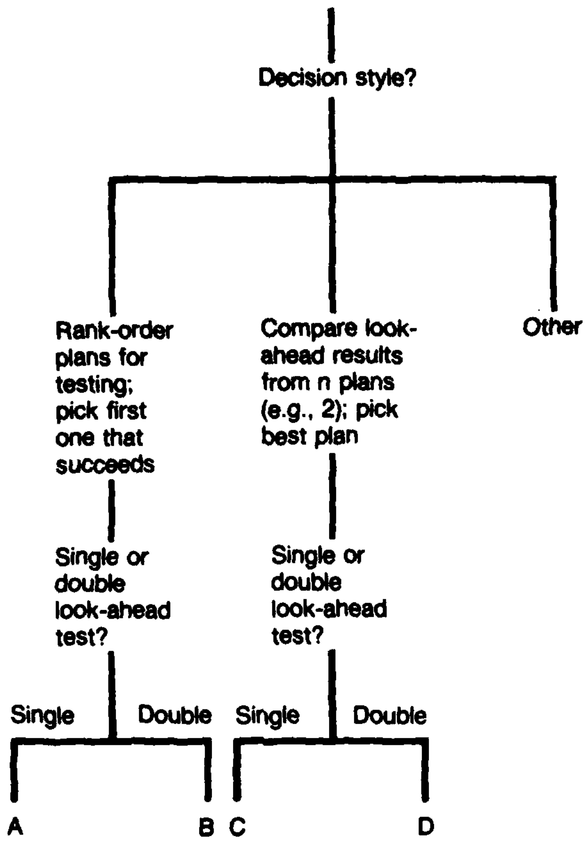

- Optional decisionmaking style, such as styles doing formal comparisons of look-ahead results with alternative plans, are permitted, but are probably to be discouraged until the effort has been made to exhaust what can be accomplished more easily with heuristic rules (which can themselves be rather sophisticated, and which can call subroutines performing projections of various types that do not require a full game within a game).
- The process of actually writing decision rules is primarily a problem of substance—of understanding the strategic issues. By virtue of the general program structure and the Rand-Abel language, "programming" is no longer the limiting factor in developing Ivans and Sams.

VI. INTERIM CONCLUSIONS AND PLANS FOR THE FUTURE

VERIFICATION AND VALIDATION

Definitions

As we observed at the beginning of the report, it is too early to talk seriously about verification and validation in the usual sense of those terms. Moreover, the standards by which programs such as the NCL models should be judged are different from those applying to a program for solving differential equations. After considerable thought on the subject, we conclude that some interesting measures of success for this class of prototype models are the following:

1. Do the programs "work?" Are they free of bugs, do they perform coherently, and do they provide appropriate explanations? Do they perform quickly? Can they be changed readily?

These questions all pertain to what we would regard as *verification*. The next issue is one of *validation*: If the programs work, do they also accomplish anything? Do the underlying models have anything to do with the real world? The issues here are:[1]

2. Do the models *usually* perform as well as experienced human control teams when used in the circumstances for which they were designed?
3. Do the models take into account the variables identified, by individuals with high-level policymaking experience, as potentially important for the types of large-scale crisis and conflict under study?
4. Is the decision style of the models comfortable to human reviewers and reasonably realistic (although not uniquely so)?
5. Does working with the models materially assist analysts in understanding and improving the logic of arguments, and in identifying the crucial assumptions that distinguish one set of views from another—i.e., do the models represent a new tool for analyzing policy-level problems with substantial, even

[1]The literature is inconsistent in its use of "verification" and "validation." Our definitions correspond to those used, for example, in Zeigler (1984a) and Shannon (1975).

dominant, qualitative components? Do the models assist the analyst in seeing scenarios very different from the standard best-estimate case, which may itself be very unlikely because uncertainties make highly diverse scenarios comparably plausible?

The reader will notice that our criteria here are much less stringent than they would be if the models were to be used autonomously. The reason for this is that *the NCL models are to be used interactively, in instances in which it will be both easy and reasonable for human analysts to override occasional model stupidities.* Indeed, one of the principal virtues of such models is that they should allow analysts and other specialists to refine and tighten their own thinking, which requires a give-and-take man-machine interaction. As experience of this sort is accumulated, the models will also improve. However, it is in the nature of the problems at issue that we should not expect the models to be comprehensive.

Our criteria are also unusual in that they do not measure validity in terms of hard "data," but instead relate model behavior to the performance of national-security specialists working on problems with large uncertainties and no real "experts," and in terms of the models' ability to clarify our thinking in interactive work as mentioned above.

In the future, as we develop more specialized NCL models for particular application studies, additional criteria will be important. For example, it will be important to characterize the range of scenarios for which the NCL models are reasonably applicable. In the meantime, however, the above criteria seem stringent enough.

Interim Conclusions

As of April 1986, we have not yet used NCL models in an applications study; instead, the work has been largely in the nature of basic research. As a result, the prototypes, although operational for demonstrations and continuing research, have not yet been exposed to thorough scrutiny. They have, however, been described in general terms to a large number of people—primarily through briefings and on-line demonstrations. And, since the prototypes were completed, we have had the opportunity to conduct several "seminar war games" focused on the types of issue addressed by the NCL models. These games consisted of very small groups talking through what the issues and decisions might be in a variety of high-level crises and conflicts involving the superpowers. Some of the people involved have had many years of experience with political-military war gaming and/or policymaking. We have also had occasion to observe deliberations in several government political-military war games.

Although our conclusions on this are inherently and unabashedly subjective, we are now convinced that the prototype models go far toward capturing the issues addressed in those human war games. Since the models were intended to be nothing more than prototypes, it would be silly to attempt true "validation." Nonetheless, it seems that the approach we have taken is valid and that it should be possible to enrich the NCL models enough to meet the criteria expressed above. Furthermore, we would claim that the prototypes (which are, in fact, under constant change) have passed the last three criteria already—in part because human-team discussion of strategic issues such as escalation in crisis is much less complex and sophisticated than is often assumed (see also Davis and Stan, 1984).

Continuing in our evaluation, the programs have not really settled down enough to allow rigorous verification, but they have been exercised extensively enough so that—with respect to criterion (1)—we can claim that the programs work, perform coherently, provide adequate explanations, perform very quickly, and can be changed readily.

To put the matter differently, then, we consider the prototype *development to have* been highly successful. At this point, we are no longer limited by modeling technique or by the complications of communicating with recalcitrant and literal-minded machines: The limiting factor is now rigorous research and analysis about national behaviors.

PLANNING THE NEXT STEPS

Model Improvements

The most important next step for our research is the straightforward one of enriching the models enough so that they can be used in serious studies and automated war games without much human intervention. We have recently begun the enrichment process and are optimistic about what can be accomplished over the next few years. Nonetheless, we should not understate the magnitude of the problems involved. Currently, we are aware of the need to enrich the models in at least the following respects:

- Enhance sensitivity of decision rules to perceptions of the opponent, third countries, and laws of war.
- Enhance sensitivity to the time dimension (e.g., the pace of events and its relationship to rationality).

- Enhance sensitivity to command and control effects generally, which will be increasingly feasible as other RSAC models dealing with command and control emerge from parallel research, and as the RSAC develops increasingly sophisticated analytic war plans.
- Broaden the class of scenarios and detailed situations for which the rule-sets are potentially applicable (e.g., conflicts focused on the Far East or Middle East).
- Develop a reasonable *set* of Ivans and Sams representing different grand strategies and temperaments. Explore in some depth the issue of U.S.-Soviet asymmetries.
- Improve the quality of decisionmaking by performing multiple look-aheads (e.g., testing a plan with both best-estimate and worst-case assumptions about the opponent's likely behavior).
- Improve the quality of decisionmaking by making explicit plan comparisons on the basis of look-aheads or more limited projections.
- Add optional stochastic features to the most critical of decision points.

As work in these areas proceeds, it will be both possible and interesting to perform more rigorous experiments comparing model decisions with those of humans and observing the degree to which human teams will make different decisions if prompted by the suggestions of models, accompanied by logical explanations.

Applications to Other Problem Areas

Finally, we should add that much of our work should have direct analogues in other domains of policy analysis. Almost all large corporations have strategic planning functions, as do government agencies. Our work represents an unusual, and probably unique, effort to combine in a large-scale complex problem area the techniques of both rule-based heuristic modeling and traditional time-oriented simulation of processes and events, and to do so in a game-structured paradigm. It seems likely that similar efforts would prove useful in other domains in which one sees adversarial processes (or even the dynamic response of the "environment," which could include government tax policy) and a combination of organizational (cybernetic) and strategic (roughly, rational-analytic, but not necessarily utility-oriented) behaviors. We hope to explore some of these issues in future work.

BIBLIOGRAPHY

Abelson, R. P., "The Structure of Belief Systems," R. C. Schank and K. M. Colby (eds.), *Computer Models of Thought and Language*, Freeman, San Francisco, 1973.

Allison, Graham T., *Essence of Decision*, Little, Brown & Co., Inc., Boston, 1971.

Anderson, John R., *Cognitive Psychology and Its Implications*, W.H. Freeman and Co., San Francisco, 1980.

Arrow, Kenneth J., "Utilities, Attitudes, Choices: A Review Note," in *Collected Papers of Kenneth J. Arrow, Vol. III: Individual Choice Under Certainty and Uncertainty*, Harvard University Press, Cambridge, Massachusetts, 1984.

Axelrod, Robert (ed.), *Structure of Decision: The Cognitive Maps of Political Elites*, Princeton University Press, Princeton, New Jersey, 1976.

Brams, Steven J., *Superpower Games: Applying Game Theory to Superpower Conflict*, Yale University Press, New Haven, Connecticut, 1985.

Bundy, William P. (ed.), *The Nuclear Controversy* (a book of reprinted articles from *Foreign Affairs*), Meridian Books, Inc., New York, 1985.

Carbonell, J. G., Jr., "POLITICS: Automated Ideological Reasoning," *Cognitive Science*, Vol. 2, 1978.

Charniak, Eugene, and Drew McDermott, *Introduction to Artificial Intelligence*, Addison-Wesley Publishing Company, Inc., Reading, Massachusetts, 1985.

Cyert, Richard M., and James G. March, *A Behavioral Theory of the Firm*, Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1963.

Davis, Paul K., *Rand's Experience in Applying Artificial Intelligence Techniques to Strategic-Level Military-Political War Gaming*, The Rand Corporation, P-6977, April 1984. (Also in the *Proceedings of the Summer Computer Simulation Conference of the Society for Computer Simulation, July 23-26, 1984*, Boston, Massachusetts, 1984.)

Davis, Paul K., *Game-Structured Analysis as a Framework for Defense Planning*, The Rand Corporation, P-7051, January 1985a.

Davis, Paul K., *Applying Artificial Intelligence Techniques to Strategic-Level Gaming and Simulation*, The Rand Corporation, P-7120, November 1985b. (Also in Maurice S. Elzas (ed.), "Knowledge-

Based Modelling and Simulation Methodologies 3rd Working Conference," North-Holland Publishing Company, Amsterdam, forthcoming.)

Davis, Paul K., "A New Analytic Technique for the Study of Deterrence, Escalation Control, and War Termination," in *Artificial Intelligence and National Security*, to be published by Lexington Books, D. C. Heath and Company, Lexington, Massachusetts. (Also available as P-7224, The Rand Corporation, May 1986.)

Davis, Paul K., and Cindy Williams, *Improving the Military Content of Strategy Analysis Using Automated War Games—A Technical Approach and an Agenda for Research*, The Rand Corporation, N-1894-DNA, June 1982.

Davis, Paul K., Peter J. E. Stan, and Bruce W. Bennett, *Automated War Gaming as a Technique for Exploring Strategic Command and Control Issues*, The Rand Corporation, N-2044-NA, November 1983. (Also in the *Proceedings of the 6th MIT/ONR Workshop on C3 Systems*, Massachusetts Institute of Technology, Cambridge, 1983.)

Davis, Paul K., and Peter J. E. Stan, *Concepts and Models of Escalation*, The Rand Corporation, R-3235, May 1984.

Davis, Paul K., and James A. Winnefeld, *The Rand Strategy Assessment Center: An Overview and Interim Conclusions about Utility and Development Options*, The Rand Corporation, R-2945-DNA, March 1983.

Drucker, Peter F., *Management: Tasks, Responsibilities, Practices*, Harper & Row, New York, 1974.

Fain, J., F. Hayes-Roth, H. Sowizral, and D. Waterman, *Programming in ROSIE: An Introduction by Means of Examples*, The Rand Corporation, N-1646-ARPA, February 1982.

Forrester, Jay W., *Urban Dynamics*, The M.I.T. Press, Cambridge, Massachusetts, 1969.

George, Alexander L., *Presidential Decisionmaking in Foreign Policy*, Westview Press, Boulder, Colorado, 1980.

Glaser, Charles, and Paul K. Davis, *Treatment of Escalation in the Rand Strategy Assessment Center*, The Rand Corporation, N-1969-DNA, April 1983.

Goldberg, A., and Alan Kay, *Smalltalk-72 Instruction Manual*, SSL 76-6, Xerox PARC, 1976.

Hayes-Roth, Frederick, Donald A. Waterman, and Douglas B. Lenat (eds.), *Building Expert Systems*, Addison-Wesley Publishing Company Inc., Reading, Massachusetts, 1983.

Hemsley, John, *Soviet Troop Control*, Brassey's Publishers Limited, New York, 1982.

Jacobsen, Mark, Robert Levine, and William Schwabe, *Contingency Plans for War in Western Europe, 1920–1940*, The Rand Corporation, R-3281-NA, June 1985.

Janis, Irving L., and Leon Mann, *Decision Making*, The Free Press, New York, 1977.

Jones, William M., *Modeling Soviet Behavior and Deterrence: A Procedure for Evaluating Military Forces*, The Rand Corporation, R-1065-PR, June 1974.

Kahan, James P., *How Psychologists Talk about Risk*, The Rand Corporation, P-6403, October 1979.

Kahan, James P., and Amnon Rapoport, *Theories of Coalition Formation*, Lawrence Erlbaum Associates, Hillsdale, 1984.

Kahan, James, William Jones, and Richard Darilek, *A Design for War Prevention Games*, The Rand Corporation, N-2285-RC, May 1985.

Kahan, James P., William L. Schwabe, and Paul K. Davis, *Characterizing the Temperaments of Red and Blue Agents—Models of Soviet and U.S. Decisionmakers*, The Rand Corporation, N-2350-NA, September 1985.

Kahn, Herman, *On Escalation: Metaphors and Scenarios* (rev. ed.), Penguin Books, Inc., Baltimore, Maryland, 1965.

Kahneman, Daniel, Paul Slovic, and Amos Tversky (eds.), *Judgment Under Uncertainty: Heuristics and Biases*, Cambridge Press, Cambridge, 1982.

Keeney, Ralph L., and Howard Raiffa, *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*, John Wiley & Sons, Inc., New York, 1976.

Knuth, Donald, *The Art of Computer Programming*, 2d ed., Addison-Wesley Publishing Company, Inc., Reading, Massachusetts, 1973.

Lefebvre, Vladimir, *Algebra of Conscience*, Reidel Press, Boston, Massachusetts, 1982.

Lefebvre, Vladimir A., and Victorina D. Lefebvre, *Reflexive Control*, Vol. I, Science Applications, Inc., SAI-84-024-FSRC-E, February 1984.

Lefebvre, Vladimir A., and Victorina D. Lefebvre, *Reflexive Control*, Vol. II, Science Applications, Inc., SAIC-85/6016&FSRC/E, 1985.

Luce, Duncan R., and Howard Raiffa, *Games and Decisions: Introduction and Critical Survey*, John Wiley & Sons, Inc., New York, 1957.

McArthur, D., and Philip Klahr, *The ROSS Language Manual*, The Rand Corporation, N-1854-AF, September 1982.

Newell, Allen, and H. A. Simon, "Heuristic Problem Solving: The Next Advance," *Operations Research*, Vol. 6, 1958.

Owen, G., *Game Theory*, 2d ed., Academic Press, Inc., New York, 1982.

Raiffa, Howard, *Decision Analysis*, 2d ed., Addison-Wesley Publishing Company Inc., Reading, Massachusetts, 1970.

Randers, Jorgen, *Elements of the System Dynamics Method*, The M.I.T. Press, Cambridge, Massachusetts, 1980.

Schank, R. C., and K. M. Colby (eds.), *Computer Models of Thought and Language*, W. H. Freeman and Co., San Francisco, 1973.

Schank, Roger C., with Peter G. Childers, *The Cognitive Computer: On Language, Learning, and Artificial Intelligence*, Addison-Wesley Publishing Company Inc., Reading, Massachusetts, 1984.

Schelling, T. C., *Arms and Influence*, Yale University Press, New Haven, Connecticut, 1966.

Schwabe, William, and Lewis M. Jamison, *A Rule-Based Policy-Level Model of Nonsuperpower Behavior in Strategic Conflicts*, The Rand Corporation, R-2962-DNA, December 1982.

Shannon, Robert E., *Systems Simulation: The Art and Science*, Prentice-Hall Inc., Englewood Cliffs, New Jersey, 1975.

Shapiro, Norman Z., H. Edward Hall, Robert H. Anderson, and Mark LaCasse, *The RAND-ABEL™ Programming Language: History, Rationale, and Design*, The Rand Corporation, R-3274-NA, August 1985a.

Shapiro, Norman Z., H. Edward Hall, Robert H. Anderson, and M. LaCasse, *The RAND-ABEL™ Programming Language: Reference Manual*, The Rand Corporation, N-2367-NA, October 1985b.

Shlapak, David, William Schwabe, Mark Lorell, Yoav Ben-Horin, *The Mark III Scenario Agent: A Rule-Based Model of Third-Country Behavior in Superpower Crises and Conflict*, The Rand Corporation, N-2363-NA, October 1985.

Shubik, M., *Game Theory in the Social Sciences: Concepts and Solutions*, The M.I.T. Press, Cambridge, Massachusetts, 1982.

Shukiar, Herbert, *Automated War Gaming: An Overview of the Rand Strategy Assessment Center*, The Rand Corporation, P-7085, May 1985.

Simon, Herbert, *Sciences of the Artificial*, 2d ed., The M.I.T. Press, Cambridge, Massachusetts, 1980.

Simon, Herbert A., *Models of Bounded Rationality: Behavioral Economics and Business Organization*, The M.I.T. Press, Cambridge, Massachusetts, 1982.

Steinbruner, John D., *The Cybernetic Theory of Decision: New Dimensions of Political Analysis*, Princeton University Press, Princeton, New Jersey, 1974.

Von Neumann, J., and Oskar Morgenstern, *Theory of Games and Economic Behavior*, 3d ed., Princeton University Press, Princeton, New Jersey, 1953.

Zeigler, Bernard P., *Theory of Modelling and Simulation*, Krieger Publishing Co., Malabar, Florida, 1984a.

Zeigler, Bernard P., *Multifaceted Modelling and Discrete Event Simulation*, Academic Press Inc., New York, 1984b.

ATE

LMED

-8