

A RAND NOTE

**Some Lessons Learned from Building Red Agents
in the RAND Strategy Assessment System (RSAS)**

Paul K. Davis

December 1989

RAND

The research described in this report was sponsored by RAND's National Defense Research Institute, a federally funded research and development center supported by the Office of the Secretary of Defense, Contract No. MDA903-85-C-0030.

The RAND Publication Series: The Report is the principal publication documenting and transmitting RAND's major research findings and final research results. The RAND Note reports other outputs of sponsored research for general distribution. Publications of The RAND Corporation do not necessarily reflect the opinions or policies of the sponsors of RAND research.

A RAND NOTE

N-3003-OSD

**Some Lessons Learned from Building Red Agents
in the RAND Strategy Assessment System (RSAS)**

Paul K. Davis

December 1989

**Prepared for the
Office of the Secretary of Defense**

**A research publication from
The RAND Strategy Assessment Center**

RAND

PREFACE

This Note contains the text of an oral presentation delivered to the third "Thinking Red in War Games" workshop held at the National Defense University in June 1988. The work underlying the presentation was accomplished in the RAND Strategy Assessment Center (RSAC), which the author directs. The RSAC is part of RAND's National Defense Research Institute (NDRI), a federally funded research and development center sponsored by the Office of the Secretary of Defense.

SUMMARY

This talk describes some of the lessons learned from developing political-level and theater-commander-level decision models for use in the RAND Strategy Assessment System (RSAS). These knowledge-based models (called Red Agents) can represent alternative mindsets and alternative military strategies. The act of building them can greatly improve insights about how to “think Red,” in part because it encourages going beyond cliches and simple-minded versions of alleged Soviet doctrine. Once developed, the models can be used in both gaming and simulation. They provide Blue players and analysts with a range of competent, coherent, Red-oriented antagonists. The models are also concrete representations of knowledge that can be passed from one study to another, or from one generation of analysts to another. It is now routine in RAND work using the RSAS to start a new study by reviewing the available Red Agent strategies and adapting them to the new study’s special needs. Because the models provide a coherent starting point, a considerably greater part of the effort can be on the special needs and research.

I begin with a short summary of RSAS architecture and a discussion of why one might want to have Red Agents. I then illustrate the use of Red Agents with two examples, one dealing with the assessment of NATO’s flexible-response strategy as a suitable mechanism for “reestablishing deterrence” in the event of war in Europe, and the other dealing with the representation of alternative theater-level Red operational strategies for force employment in the Central Region. The conclusions are that working with Red Agent models has proven quite valuable already and has enormous, even revolutionary, potential. The assessment of flexible response illustrates how the rigor required in modeling can change one’s mind fundamentally about escalation and de-escalation processes. The second application reminds us how critical the strategy variable is (or should be) in analysis, and demonstrates that it can now be highlighted and explored, rather than left buried in the script-like data bases that drive most combat simulations.

CONTENTS

PREFACE.....iii
SUMMARY.....v
Section
I. INTRODUCTION..... 1
II. THINKING ABOUT FLEXIBLE RESPONSE WITH RED NATIONAL
COMMAND LEVEL MODELS 7
III. USING RED AGENTS TO THINK ABOUT OPERATIONAL
STRATEGIES..... 15
BIBLIOGRAPHY.....21

I. INTRODUCTION

My purpose today is to describe some of what my colleagues and I have learned from developing and using models of Soviet military and political behavior, models called Red Agents. Our work has been part of a major long-term effort to combine some of the better features of human war gaming and analytic modeling in something called the RAND Strategy Assessment System (RSAS), which has the following major features:¹

- It is a system for military and political *analytic war gaming*.
- It has optional decision models called Red, Blue, and Green Agents, in addition to a suite of combat models called, collectively, Force Agent or CAMPAIGN.
- It can be used in several modes: (a) closed simulation with agent pitted against agent; (b) hybrid operations with humans competing against agents; or (c) human gaming in which the models provide bookkeeping and may be used by the players as a source of building-block orders (e.g., the orders for the deployment of Reforger forces) and may even be used as decision aids (e.g., as when players run closed simulations to test alternative strategies before making their own decisions).
- It has a strategic-level perspective (i.e., theater commander or national), with a mix of theater-level and operational-level resolution.
- It is global in scope, so that users can consider a single theater or multiple theaters.
- It covers conventional, theater-nuclear, and intercontinental-nuclear combat.

All in all, it is a highly *integrated* system designed from the top down.

The Red and Blue Agents are, of course, rather complex. Focusing on Red here, Fig. 1 shows the Red Agent to be a *hierarchy* of models. The Defense Council is the political level, deciding issues such as whether and when to go to war and for what purposes, whether to escalate, and whether to terminate. The Supreme High Command receives a statement of strategy from the Defense Council and translates that into the choice of a war plan and particular instructions within that plan (e.g., the preparation time permitted and the authorization to overfly or cross through neutral countries). This global plan actually has components for each of the military theaters worldwide (e.g., the High Command of Forces, West, which would direct war in what we call Europe's Central Region).

¹See also Davis, Bennett, and Schwabe (1988), Davis and Hall (1988), and Bennett, Jones, Bullock, and Davis (1988).

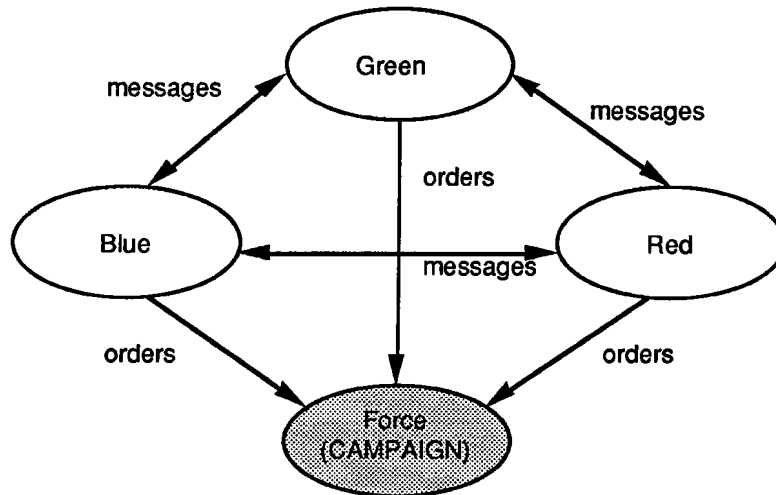


Fig. 1a—Top-level RSAS structure (simplified)

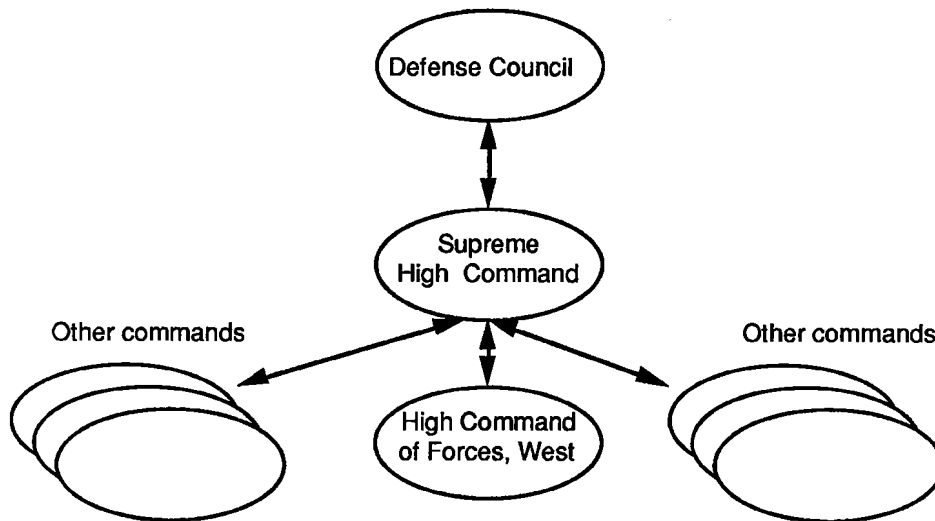


Fig. 1b—Structure of the Red Agent (simplified)

There is a fundamental difference, however, between these models and real-world plans: the models must include instructions for the computer about how the plan is to be adapted as circumstances unfold. That is, the decision models composing the Red Agent must represent not

only prewar war plans, but also the adaptations that commanders might make in the course of the war.

Why do we even want a Red Agent? Some of the principal reasons are as follows:²

- *Analytic war gaming requires* controlling decision variables as well as technical variables, since decisions determine outcomes.
- Red Agents can be repositories of knowledge (e.g., a Red Agent developed for one study can be used in successive studies so that the knowledge gained is not lost as the participants in the first study rotate out to new assignments).
- Blue players and analysts can test new concepts of operations against Red Agents without themselves being Red specialists (although more stringent testing would require human Red play).
- Building Red Agents provides structure and context for “thinking Red” seriously.

A standard question at this point is how we validate Red Agent models. The answer is that there can be no real validation without having wars. Instead of asking for true validation, what we should be asking is the following (Davis, Bankes, and Kahan, 1986):

- Does building Red Agents sharpen thinking and add insights?
- Can knowledge be passed on by using Red Agents?
- Do Red Agents focus on the “right” variables and issues as judged by experts?
- Do Red Agents perform well *most* of the time, so that with moderate human oversight they are useful in games and simulations?
- Can working with Red Agents add anything truly new and exciting?

We now have considerable experience, and the answers are, respectively: Yes, Yes, Some do and some do not, Some do and some do not, and Yes. In other words, the technique is excellent, but execution depends on who is doing the work. It has always been that way with models, and so will it always. There is an important point here, however: Red Agent models are not expected to be “correct” because we do not even know precisely what that means. They are instead supposed to be useful *tools* for understanding Red and for testing our own concepts and

²See also Davis and Schwabe (1985), a paper developed for the first “Thinking Red” workshop.

capabilities. We need many different Red Agents because there are many possible Red strategies, and even matters such as doctrine are not really so clear-cut when one looks into them closely.

In what follows I will describe two applications of our Red Agents. The first uses decision models to address NATO's concept of flexible response. The second describes the use of military command-level models representing alternative military strategies.

II. THINKING ABOUT FLEXIBLE RESPONSE WITH RED NATIONAL-COMMAND-LEVEL MODELS

Late last year I conducted some informal experiments with special variants of our Red national-command-level models (NCL models), the Red Agent version of which corresponds to the Defense Council. The purpose was to think more deeply than usual about whether flexible-response theory makes sense. Would it really work if applied in crisis or conflict? The approach was to build some intendedly realistic “Ivans” and observe their responses to Blue escalation in simulations, to then think about those responses, and finally to draw some conclusions—carefully. Note that the models were to be used more as aids to thinking than as calculational aids, much less as answer machines. The question was whether by using the models I would learn something, just as one often learns something from human political-military gaming even though it is not rigorous. As to how one decides whether the insights gained are correct, the answer is that one must judge that in the same way we judge the validity of all the other insights we gain from our real-world experiences, war games, reading of history, and conversations.³

Flexible-response strategy, you will recall, is at the heart of NATO’s deterrent concept. The idea is that NATO will be prepared to attempt a conventional defense if invaded. However, if that defense falters, NATO will be prepared to escalate to the use of nuclear weapons—perhaps by launching a very small “demonstrative” strike against military targets—to convince the Warsaw Pact that it should cease and reverse its aggression. The phrase often used here is that NATO will be prepared to use nuclear weapons first, if necessary, to reestablish deterrence. And, if such an initial use fails, it will be prepared to go farther. Indeed, a basic element of the strategy is a “linkage” to the U.S. strategic nuclear arsenal and prospects for general intercontinental nuclear war. There were strong arguments in the development of the strategy during the 1960s, and the debates continue. The question is whether we can think about it insightfully by using models.⁴

As background here, I would point out the following. First, technical studies do not even address the issue, instead focusing on such matters as the fraction of targets that can be covered by NATO’s tactical nuclear weapons—even though that issue seems not to have been deemed especially important to the framers of flexible response. Second,

³For an excellent discussion of this see Schelling (1987), an article of interest to anyone concerned with higher-level war gaming.

⁴Early thoughts on this matter are discussed in Davis and Stan (1984).

human games that include nuclear escalation vary enormously in quality and seldom attempt to treat the details of flexible response realistically because the games usually have a purpose that *requires* the escalation to occur. In some of the exceptional games in which escalation decisions are debated seriously, the character of the games is often colored strongly by assertions about what Soviet military doctrine would require the Soviets to do, even though the Soviet leaders making the decisions might not be especially interested in what U.S. military people believed they were required by U.S. interpretations of Soviet doctrine to do. There have been many games over the years in which the Soviet response to limited NATO nuclear use was a large-scale attack, often an attack on the United States as well as Western Europe. In other games the issue of Soviet response has not arisen, because the Blue players simply do not want to use nuclear weapons. And so it goes. There has been little convergence of views as the result of human war gaming on these matters. And, even if some would disagree with this because they felt that *they* had learned deep insights, those lessons have not been passed on to others very well, because the *details* of the more serious war games are seldom documented and distributed.

One final point as background. It is my strong impression that the prevailing view among defense intellectuals in the United States and Western Europe is that NATO's policy should be against *early* first use of nuclear weapons. NATO does not have strong enough conventional defenses to forgo the nuclear deterrent, but at least, in this view, NATO should plan on an initial conventional defense.

So much for background. To explore this issue I conducted perhaps twenty or thirty simulations, in each of which the Blue model (called "Sam," in an anthropomorphic manner that adds color but admittedly causes some trouble as well) tried dutifully to conduct "himself" according to flexible-response theory as I understand it. It is important to understand the nature of the particular Red model I used in this work, because we have other types of Red model that would also be plausible. The Red model I used was opportunistic, but also rational and even pragmatic. Indeed, he would not even consider invading Europe if he really believed that NATO would use nuclear weapons.⁵ The basic scenario in all of the simulations was one in which Ivan started the war with the belief that it would be accomplished quickly and conventionally. The

⁵Here and in what follows I have taken literary license in describing the models' "thinking." Actually, the models consist of nothing more than a set of If...Then rules, plus lengthy commentary by the model builders trying to explain the mindset and behavior patterns they were trying to achieve with a particular, intendedly self-consistent, set of rules. When the models run there are explanation logs, but those are much less picturesque than my language here.

strategy involved a short-mobilization attack which, in favorable cases, permitted a Pact victory within 2-3 weeks if war remained conventional. Why did Ivan think the war would remain conventional? Remembering that this is all rationale underlying the model, not part of the model itself, the likely reason is that the Blue President and his NATO allies had in recent years become No-first-users—perhaps even running for office under such a label. And, probably, NATO had let its nuclear modernization slip and the intensity of its nuclear training drop. More formally, the model of Ivan has parameter settings that determine whether, at the outset, he sees opportunities, what type of opponent he believes he is facing, and so on.

The simulations, then, dealt with a very special class of scenarios in which Ivan had made a very, very serious mistake. He thought Sam would not use nuclear weapons, but Sam in fact did use them. I happen to believe this is a *relatively* plausible case, because I do not believe the Soviets would invade anywhere if they expected nuclear war as a consequence.

The simulations varied in a number of dimensions. In some cases I assumed as an exogenous constraint that NATO could not reach a nuclear decision “early,” regardless of merit. In other cases, I assumed that Sam would escalate as early as he concluded that was a wise thing to do. In other words, the time of earliest-first-use was a parameter to be varied. I also varied some of the technical assumptions about relative force effectiveness and the details of scenario (e.g., the exact number of mobilization days)—primarily to get a range of wars.

Rather surprisingly, it did not take very much tinkering with input assumptions to change drastically the nature of what happened. In all cases, Sam eventually conducted a small “demonstrative” nuclear attack in an attempt to reestablish deterrence. The various games developed in several different ways, which can be grouped as follows:

- A: Ivan conducted a broad but limited retaliatory nuclear attack in Europe; NATO responded with a broader attack and Ivan escalated to intercontinental nuclear warfare.
- B: Ivan agreed to war termination with the expectation of no net gains from the war.
- C: Ivan responded in kind with a very few nuclear weapons striking tactical targets.
- D: Ivan “swallowed” the nuclear strike and pressed on with the campaign.

One of the satisfying features of all this was that the same Ivan exhibited such strikingly different behaviors as a function of context at the time of nuclear first use. This had been a goal of the original model-building effort, because—at the urging of RAND Soviet-specialist colleagues—I had sought to avoid making the Ivans and Sams into simple-minded stereotypes

(Davis and Stan, 1984, and Davis, Bankes, and Kahan, 1986). Nonetheless, why were the outcomes so drastically different? An outside observer might even have called them “random.”

In one case, the game outcome resulted from what might be considered a bug or from what might be considered a very realistic manifestation of political-military misunderstanding. The colleague who had developed Ivan’s package of nuclear options interpreted the requirement for a “very limited nuclear option” to be something larger and more effective than what Sam’s political-level logic “had in mind.” As a result, when Sam launched his demonstrative strike, Ivan concluded that he should respond in kind—but that produced something a good deal larger than Sam was expecting. When Sam saw the response, he concluded that limited use had failed and that Ivan was a nuclear warfighter, not a bargainer. He then granted authority for his own commanders to use tactical nuclear weapons and the war eventually escalated from there—not surprisingly—to intercontinental nuclear war, because Ivan did not really believe war could be contained at this point. Why? In large part because he concluded that he had completely misunderstood Sam and that Sam was now intent on a war to the finish and would surely attack the Soviet Union at some point. In other words, the crucial decision in this entire chain could be said to be the decision by Ivan about the nature of his opponent. And the cause of that crucial decision could be said to be a simple misunderstanding between what political and military figures might regard as a “small” and “restrained” use of nuclear weapons.

The Class-B results are probably the ones that will be most surprising at first glance. We are accustomed to ignoring the possibility of the Soviet Union capitulating. Frankly, I would not want to bet much on this case happening if war began: war termination is notoriously difficult (although all nuclear wars to date have been terminated shortly after nuclear use). Nonetheless, I had been especially interested in seeing whether there were *any* cases in which plausible rules could lead to flexible response working in the sense of reestablishing deterrence, and there were a few such cases. If one puts aside the accumulated experience of studies ignoring termination and thinks about the issue logically, it is perhaps not hard to guess the circumstances in which Ivan quit. Or perhaps it is not so obvious. In any case, for the limited range of simulations I conducted, this particular Ivan quit under one principal circumstance: when NATO initiated first use immediately, indeed on or about D-Day. Ivan’s reasoning was, essentially, that he had completely miscalculated, that the sides were now on the threshold of general nuclear war, and—importantly—that the die was not completely cast and the prospects if war continued were exceedingly uncertain. Would a real Ivan have quit? Who knows? I would argue, however, that it is not implausible that *if* there are scenarios in which flexible response would work, the immediate-use scenarios must be prominent among them.

By contrast, in virtually all other cases, Ivan's decision about nuclear use came after war had been ongoing for some time and Sam's defenses were crumbling. In these cases, Ivan was still distraught about the nuclear use, but he considered the die cast and noted that victory was within his grasp. He also "knew" that anything that NATO did could be countered in kind. Why should he quit at this stage? Indeed, even though Ivan had to change his model of Sam at this point, the new model of Sam was one of a nearly defeated but rational adversary trying one last gamble before capitulating. This was very different from the situation in which NATO used nuclear weapons on D-Day, thereby forcing Ivan to completely restructure his thinking about NATO's (or Sam's) resolve and "rationality."

Now, I must emphasize that whatever Sam and Ivan did they did as the result of rules that I or my colleagues put into them. Thus, I could have rigged the rules to produce precisely what happened—or something quite different. The fact is, however, that the games frequently did not play out the way I had expected when I was merely trying to put together logical self-contained representations of particular behavior patterns for the United States and Soviet Union. As so often happens in simulation, exercising the model revealed things that were there all the time but that had not been very clear. The most important observation seems to be that *if* NATO wants to depend on flexible response, then it should most certainly have options ready for immediate first use rather than accepting the concept that "no early first use" is automatically virtuous. Indeed, after thinking rather deeply about all this subsequently, I am now convinced (not because "the models said so," but because of the thinking they caused me to do) that *the basic flow of mainline no-early-first-use thinking is flat-out wrong and very likely to fail if used in time of conflict*. At the same time, I also believe that conventional defense is feasible with moderate effort by NATO, and I would vastly prefer not to depend on the nuclear deterrent except in the most general way.

The next lesson was that other people, who had had nothing to do with the original model building, learned from observing the simulations. For example, some of my colleagues were initially incredulous that Ivan would ever terminate (they had never really thought much about termination). On being pinned down about what better rules to write for Ivan, however, they found themselves taking more seriously the chain of logic exhibited: that the Soviets would be most likely to invade Europe if they were convinced there would be no nuclear war, and if they then found themselves very early in precisely that nuclear war that was not supposed to happen, then...maybe, just maybe...they would throw in the towel. While none of us would bet on nuclear war being containable, the possibility of a satisfactory termination in this way seemed less implausible than in other ways.

Similarly, some of my colleagues were displeased when Ivan did not respond to Sam's first use with a much larger nuclear use (perhaps they had in mind claims about Soviet military

doctrine). Upon reviewing the logic, however, and noting that Ivan's decision *not* to escalate further led to his victory, they changed their minds. I should elaborate further on this. In Cases C and D, Ivan typically won the war—i.e., the Soviets gained their war objectives in terms of changing the map of Europe. In none of these other simulations did Sam continue to escalate. In these cases, at least, the demonstrative nuclear use was the last nuclear use. Now, in some cases *Sam* was willing to continue and expand limited nuclear use (I noted that this particular Sam was programmed to be a real believer in NATO's flexible-response strategy), but some of the NATO allies as represented by the RSAS' Green Agent (and the particular parameter values used for it in these simulations), upon seeing that the demonstrative use of nuclear weapons had failed *and* that the war was all but lost militarily, concluded that the game was over. They did not approve further nuclear use, and Sam was not about to escalate further without the support of his allies.⁶

Another lesson, then, is that the success of flexible response would depend on the continued cohesion of key alliance members and that proponents of flexible response should not assume that such cohesion could be maintained, even if there were agreement on the initial use of nuclear weapons. Moreover, waiting until the war was being lost militarily would mean that alliance resolve would indeed be questionable if the demonstrative use brought back a nuclear response (unless, of course, the nuclear use also had a profound military effect that could not be countered by a response in kind, something that was seldom the case in these simulations).

Before shifting to the second example of Red Agent analysis, let me comment briefly on some of the difficulties in writing serious political-level Red Agents. One of the principal obstacles is that the majority of "Red specialists" are steeped in military doctrine, not policy-level thinking. Second, many specialists over the years have misrepresented Soviet military doctrine. To be blunt, Sokolovsky never said what they said he said. When one actually looks up the various quotations that allegedly prove Soviet propensities to behave in various ways, one finds them filled with words like "could" and "might." Moreover, a *universal* tenet of military doctrine is that the proper action at the time depends on the circumstances, and it is certainly a tenet of Soviet thinking in particular. Objective facts, not subjective preconceptions, must dominate decisions. Calculate, calculate, calculate! Yet, if we were to build top-level Red Agents following the advice of some Red Team folks, their behavior would be irrational—leading to the destruction of the Soviet Union. Now, it may well be that nuclear war would be uncontainable, for reasons I have elaborated upon at length elsewhere (Davis, 1989), but we should not assume

⁶Although the particular simulations did not produce such results, the underlying rules in the Ivan and Sam models allow for intercontinental nuclear war to be catalyzed in some instances by the independent nuclear use of French or British strategic forces, or by "accidental" escalations.

that rational Soviet leaders would make foolish decisions because of their doctrine allegedly requiring them to do those foolish things.

I was particularly impressed several years ago by the writings of a late Soviet specialist and RAND colleague, Nathan Leites. Leites was extremely sensitive to the nuances of language, and his discussion of Soviet views of nuclear deterrence were remarkable for their time (1982), since many experts in the United States were then claiming that the Soviets did not understand “deterrence,” nor even have a word for it. Coming from a background of physical science, it has always rubbed me wrong when someone claims that the Soviets do not “understand” things (although I have no difficulty seeing why they reach different conclusions, have different strategies, or hold different values), but it was good to read comments such as the following by someone of Leites’ background (see Leites, 1982, pp. 378-379):

As in a few other ways, the Kremlin is in this regard more *seriozny* than we.... It is perhaps just because the Soviets are so interested in the distinction between deterrence and war-fighting that they have kept silent about it. The war not being yet begun, this is the hour of deterrence: deterrence by the prospect of a maximum initial strike, of preemption, and of the none-or-all character of nuclear war. Once the war is on, the Authorities may adopt that “controlled” conduct about which the West (in a possible Soviet estimate) is now so prematurely chattering.

III. USING RED AGENTS TO THINK ABOUT OPERATIONAL STRATEGIES

What I'd like to do in the second example is talk more about the *military* analysis possible with Red Agents. This, in fact, is the principal application of the Red Agents now, and I expect it to remain so indefinitely. The problem I would like to consider here is how Red might constitute an invasion strategy for the Central Region. What would a Red strategy look like if we saw it? Well, first, there are such standard issues in analysis as: (1) mobilization time (e.g., 72 hours, 7 days, 3 weeks, or 3 months); (2) the main-thrust axes (e.g., the Belgian corps sector and 2 to 4 others); and (3) the use, if any, of strategic reserves. But beyond that, most work in studies goes typically into getting the data "right" on order of battle, weapons holdings, and so on. But is this all there is to strategy? The answer, of course, is assuredly not, but the other elements of strategy are either dealt with as they arise (as in gaming) or given short shrift: they are not often given the analytic attention they deserve.

What are some of these "other issues?" They include:

- Balancing one's own preparations against those of the opponent.
- Achieving some measure of strategic and operational surprise.
- Establishing conditions for early breakthroughs and strategic-level encirclements.
- Achieving success in a massive air operation to suppress NATO's tactical air forces, which could greatly disrupt the momentum of the ground-force attack.
- Coordinating events in other theaters (balancing preparations vs surprise), including the Northern and Southern Regions of Europe.
- Hedging against the unexpected.
- Anticipating possible NATO actions.

While this is by no means a complete list, it illustrates the complexity of a war plan. The last item is particularly important and relevant, and distinguishes gaming and real-world planning from standard modeling. In standard models and studies, there is relatively little adaptation by the sides as the war continues. In many simulations, the sides' forces are assigned by a rigid and predetermined *script*, which does not even allow for adaptations. With this as background, consider some of the features one might want to have in a family of Red Agent models (Davis, 1988):

- Options for basing D-Day on NATO's behavior (see Fig. 2).
- Options for changing main-thrust axes of advance as late as D-2:
 - to exploit weaknesses due to the late arrival of some NATO allies,
 - to counter surprise NATO reinforcement of main-thrust sectors.
- Options for premobilization training that would improve the effectiveness and/or shorten the arrival time of second-echelon forces.
- Adaptive logic for committing operational reserves.
- Options for the use of the Austrian corridor, either with or without the Southern Group of Forces.
- Options for breaking off or refocusing the air operation.

Figure 2 illustrates a simple logic for dealing with the first issue. The issue here is when D-Day should be. The more preparations Red makes, the more preparations NATO might make, depending on how quickly NATO begins to mobilize. In this figure it is assumed that the Red commander believes it would be unwise to attack any earlier than M+3, or any later than M+10,

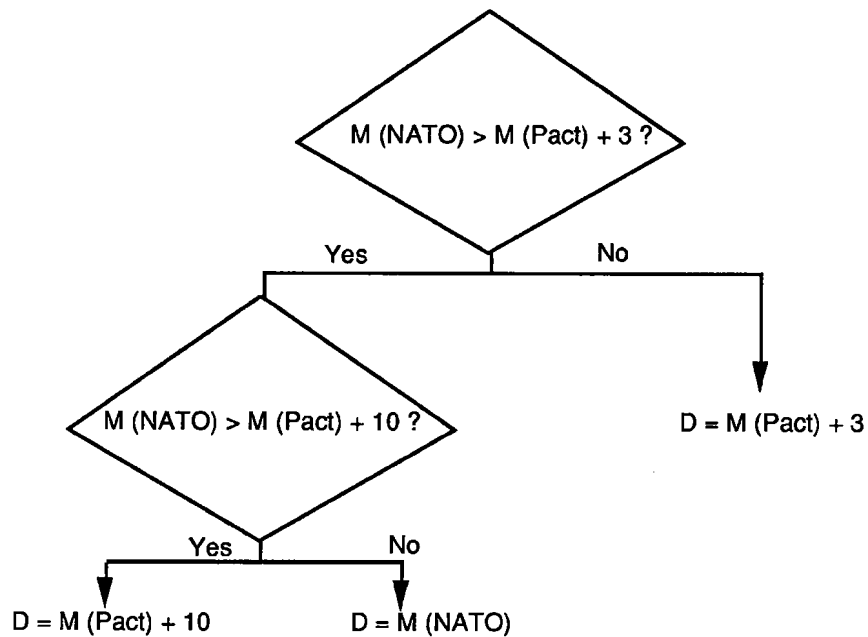


Fig. 2—Flexibility in setting D-Day as part of a strategy

but within that window of time the exact D-Day would depend on NATO's response to the Pact's mobilization. If NATO mobilized within the first three days, then the attack would be on M+3; if it mobilized sometime in the next week, the Pact would attack immediately to gain maximum benefit from NATO's initial disruption. In any case, however, D-Day would be no later than M+10.

There is nothing especially remarkable about this logic, but by having the mechanism of a Red Agent available, one is *encouraged* to write such logic down—and to think of more sophisticated plans to incorporate deception, surprise, and adaptations to NATO's actions. Human players must do some of this ad hoc during games, and if a senior officer is playing, his actions may become part of military lore for a period of time, but by and large there has been no mechanism to encourage, structure, preserve, and communicate these essential features of military strategy except in the most general way.

Again, the purpose is not to “validate” these Red Agents. Instead, the question is whether building and using Red Agents is helpful. The answer, based on considerable experience by my colleagues and myself, is an unequivocal yes. This feature of the RAND Strategy Assessment System is, in many ways, the one I have long regarded as the most important new innovation, and I have been gratified to see it come into use. Once one gets into the spirit of building Red Agents and testing them in a wide variety of scenarios, issues arise that may have been swept under the rug for years. In our work, for example, we have considered Red strategies for:

- Blunting an invasion by NATO (!), and then going on the counteroffensive.
- Dealing with unreliable non-Soviet Warsaw Pact allies.
- Dealing with horizontal escalation in other theaters.
- Responding to NATO's limited nuclear use.

The first two items are interesting in part because it has been so unusual over the years to even consider them. If we truly wish to understand how Soviet military leaders view the world and structure their plans, however, we should almost certainly be examining such cases analytically—if for no other reason than to be able to speak with authority to the Soviets about why their alleged fears of a NATO invasion are not credible, or about how—even if they fear unreliable allies—they have no need for the large numbers of forward-stationed forces in the GSFG (Group of Soviet Forces Germany). Also, I would hope that NATO's military planners would have plans to exploit problems that might arise within the Warsaw Pact alliance in the course of war.

To summarize, then, having the RSAS' analytic war plans, which are military-command-level Red Agents, forces one to consider multiple components of strategy systematically. It also encourages developing explicit adaptive logic that can be used to refine operations plans and explore the many possible stratagems of the opponent. And, finally, it provides a natural vehicle for top-down planning around objectives in the presence of constraints and uncertainty.

As of today, RAND has many skeletal analytic war plans for theaters worldwide, and a fair number of second-generation plans for the Central Region and elsewhere. I would not want to exaggerate their quality or sophistication, except by comparison with what has been used previously in studies and games. Instead, my real point is that *the technological base has been laid for a paradigm shift in the way military analysis and planning is conducted*. Today, we can see the beginnings of this shift. In ongoing studies at RAND and within the Pentagon, analytic war plans developed in last year's studies are now being used for new ones, not just by the original people, but by their successors and by people in other projects. To be sure, the plans must be adapted and updated because each study has unique interests and there are continual changes, but the plans are now providing a continuity that has been badly lacking in years past. In the past, it was common to use a static data base from a previous study in a new study, but there was much less carryover of structured *knowledge*, especially about strategy and adaptation.

One spinoff of this approach to strategic thinking is that one develops over time a library of strategies, an experience base from multiscenario simulation, and a broader perspective on the military problems we face. This, in turn, permits us to begin describing the military balance itself in a more organized way. Figure 3, for example, summarizes with a fault tree many different classes of defense failure, against all of which NATO should plan (Davis, 1988). We have been able to simulate all of these failure modes in the RSAS, and to gain a better sense for the key variables.

As a last comment on the issue of building Red (or Blue) analytic war plans, let me observe that while the methodology works and the technology it represents is a major step forward, we are dealing with only the first generation of a new class of analytic tools. As a result, working with the system is by no means as easy as it could be or will eventually be (this is a significant understatement!). Furthermore, because we are starting from a base in which there are no adaptive plans to emulate, it will take years to build up sophistication of the library of plans and plan components. This work will be neither easy nor mechanical; to the contrary, it will require military savvy, analytic talent, time, and patience. At any given time the existing plans will always have holes, some of them serious and some of them trivial but annoying. And, because expectations rise very

quickly, even faster than capabilities, I do not expect to see mature capabilities in this area for many years. Nonetheless, our initial experience has been strongly encouraging.




Fig. 3—A fault-tree depiction of potential NATO problems

BIBLIOGRAPHY

- Bennett, Bruce Wm. , Carl M. Jones, Arthur M. Bullock, and Paul K. Davis, *Main Theater Warfare Modeling in the RAND Strategy Assessment System (3.0)*, The RAND Corporation, N-2743-NA, 1988. High-level documentation on RSAS combat models.
- Davis, Paul K., Bruce Wm. Bennett, and William L. Schwabe, "Analytic War Gaming With Ivan and Sam," *Signal*, July, 1988. A short overview of the RSAS written for a relatively wide military audience.
- Davis, Paul K., *Applying Artificial Intelligence Techniques to Strategic-Level Gaming and Simulation*, The RAND Corporation, N-2752-RC, 1988. A technical paper describing to simulation and computer-science experts how various A.I. concepts have been adapted and extended in RSAS work. Published initially as a chapter in Maurice Elzas et al, (Eds.), *Simulation in the Artificial Intelligence Era*, North-Holland, Amsterdam, 1986.
- Davis, Paul K., *Studying First-Strike Stability With Knowledge-Based Models of Human Decisionmaking*, The RAND Corporation, R-3689-CC, April 1989.
- Davis, Paul K. and H. Edward Hall, *Overview of RSAS System Software*, The RAND Corporation, N-2755-NA, 1988. A relatively detailed annotated briefing providing a top-down description of the RSAS' technical architecture of competing agents.
- Davis, Paul K., *The Role of Uncertainty in Assessing the NATO-Pact Central-Region Balance*, The RAND Corporation, N-2839-RC, December 1988. Describes the author's general philosophy toward balance assessments and using such assessments for problem solving.
- Davis, Paul K., Steven C. Bankes, and James P. Kahan, *A New Methodology for Modeling National Command Level Decisionmaking in War Games and Simulations*, The RAND Corporation, R-3290-NA, 1986. Documentation of the general approach taken on national-command-level modeling in the RSAS.
- Davis, Paul K., and William L. Schwabe, *Search for a Red Agent to be Used in War Games and Simulations*, The RAND Corporation, P-7107, 1985. Discusses early intentions and treatment of Red-Blue asymmetries.
- Davis, Paul K. and Peter J.E. Stan, *Concepts and Models of Escalation*, The RAND Corporation, R-3235, 1984. Compares U.S. and Soviet concepts of escalation and sketches out an approach to modeling related decisionmaking. Argued against the then canonical image of Soviet nuclear decisionmaking by observing that military doctrine tends to deal with an "ideal war," whereas actual decisions would depend on details of context.
- Leites, Nathan, *The Soviet Style in War*, The RAND Corporation, R-2615-NA, 1982 (also available commercially). A unique description of Soviet military thinking based on the nuances of Soviet military literature.
- Schelling, Thomas, "The Role of War Games and Exercises," in *Managing Nuclear Operations*, edited by Ashton B. Carter, John D. Steinbruner, and Charles A. Zraket, The Brookings Institution, 1987. An excellent and realistic assessment.

