MEMORANDUM
RM-1829-1-PR
JUNE 1957

TECHNIQUES OF SYSTEMS ANALYSIS

H. Kahn and I. Mann

PREPARED FOR:
UNITED STATES AIR FORCE PROJECT RAND

The RAND Corporation
SANTA MONICA • CALIFORNIA

MEMORANDUM
RM-1829-1-PR
JUNE 1957

TECHNIQUES OF SYSTEMS ANALYSIS

H. Kahn and I. Mann

This research is supported by the United States Air Force under Project RAND—Contract No. F44620-67-C-0045—monitored by the Directorate of Operational Requirements and Development Plans, Deputy Chief of Staff, Research and Development, Hq USAF. Views or conclusions contained in this Memorandum should not be interpreted as representing the official opinion or policy of the United States Air Force.

DISTRIBUTION STATEMENT
Distribution of this document is unlimited.

The RAND Corporation
1700 MAIN ST • SANTA MONICA • CALIFORNIA • 90406

Published by The RAND Corporation

SPECIAL NOTE

This is the preliminary draft of a report which is being circulated for information and comment. As explained in the foreword we hope eventually to incorporate it into a book and would, therefore, appreciate any comments, criticism, ideas, and examples that readers may have. This draft began as a verbatim transcript of an informal talk and, despite some rewriting, it probably still suffers (like many such talks) from being "fashionable." We are aware that it has a number of other weaknesses and assume there are still others of which we are not aware. We hope to give it a thoughtful and leisurely review but are deferring this until we get some outside criticism.

In order to give the reader a feeling for the place this material might have in the book a table of contents of the book is given on the next page.

H. K.
I. M.

MILITARY PLANNING IN AN UNCERTAIN WORLD

CONTENTS FOR BOOK

PART ONE: TECHNIQUES OF SYSTEMS ANALYSIS[1]

---

[1] This document

[2] Has already appeared as RAND P-1165

[3] Has already appeared as RAND P-1166

[4] Has already appeared as RAND P-1167

[5] Has already appeared as RAND RM-1937

FOREWORD

This is one of three fairly independent publications which we hope to put together in a book. The three sections together are essentially a collection of talks given for various reasons, but having a common theme which will be the title of the book--Military Planning In An Uncertain World. They are all concerned with some of the contributions that Systems Analysis and Operations Research can make to this planning. The points of view expressed are personal and reflect what we have learned about planning during our association with The RAND Corporation.

This first section, as noted before, is a nearly verbatim account of a combined lecture and workshop that took one afternoon of a one week course given by The RAND Corporation to selected Air Force audiences. The course was called "An Appreciation of Systems Analysis." Its major objective was to make the students better and more critical consumers of Systems Analyses, rather than to train technicians in the art. The instructors, who are listed in the acknowledgments below, spent much of their time emphasizing that uncommon quality--Common Sense. The talk that we prepared for the course and reproduce here was mainly to illustrate techniques of Systems Analysis and was only incidentally concerned with fundamental principles. However, we did consider the principles, sometimes more than just briefly.

The second portion of the future book, Techniques of Operations Research, is to have six chapters. Five of them, "Flyaway Kits," "Programming," "Monte Carlo," "Game Theory," and "War Gaming," were prepared as talks, some for the systems analysis course and some for other occasions. The chapter "Probability and Statistics" was written to make the section more complete.

The third portion of the book, Philosophical and Methodological Comments, will start with a chapter called "Ten Common Pitfalls." This material was originally prepared as a diatribe against certain common practices. In its original form it was a somewhat immodest and immoderate document. In the interest of fair play, manners, and justice, it has been softened. In addition all classified and personal examples have been removed. These deletions unfortunately result in a slightly saccharine flavor which we deplore but cannot avoid. The second chapter, "Nine Helpful Hints,"[1] was originally written by Harvey Lynn as a satire on our "Ten Common Pitfalls." We liked it so much that we want to include it in the book. Since we have made some changes, it would be unfair on the part of the reader to blame Harvey for any specific remarks unless he verified that Harvey actually made them. The last chapter will contain miscellaneous comments. These are mainly postscripts that will probably be included in the text when we revise it. On the whole, while the comments are elementary, we think of some of them as being quite important.

We do not really know what kind of readers this report will interest. As it happens, however, we do not depend upon the market determining the audience, since we have selected every reader with some care. If you aren't interested it's our fault not yours.

The talks had a quite favorable reception by audiences as varied as laymen, physical and social scientists, military personnel, and professionals in the field, and we hope that the book will have a similar audience. We are

[1] Journal of the Operations Research Society of America, August, 1956.

not certain that they will stand up well under the bright light of print. However, if this report succeeds only in stimulating interest in either the strictly technical problems or in some of the wider public policy problems, it will have justified its existence. If, in addition, it transmits some information or pleasure, we shall be even happier.

We apologize for our prolific use of underlining. We have done it in spite of many admonitions from friends and colleagues and in full awareness that many people will find it irritating. We feel, however, that some readers will find it helpful. The alternative would call for a major revision of the talks. We are deferring this revision until we get some comments on the current version.

One last remark. The subject of this memorandum is Systems Analysis. Actually it is being written from a background of experience in Applied Mathematics as well as in Systems Analysis; a large percent of our "preachy" and "hardheaded" remarks pertain as much to the former field as to the latter and originally arose from experience in that field.

Herman Kahn

Irwin Mann

ACKNOWLEDGMENTS

A book like this is to some extent a product of an environment and reflects the work of many minds. We therefore begin by acknowledging our general debt to the whole profession of Systems Analysis and to the staff of The RAND Corporation in particular. In the latter group, Albert Wohlstetter and Andrew Marshall especially have played important roles in our intellectual development. They introduced us to many of the ideas in this report and we have had many discussions with them about these ideas. In addition they can be considered as being (in a creative sense) co-authors with us of the "Ten Common Pitfalls" chapter although they may wish and should be allowed to disown specific paragraphs.

We owe special gratitude to Leonard J. Savage of the University of Chicago and Bernard Brodie of The RAND Corporation who made very detailed criticisms of almost every idea in the report.

We have specific debts to the other lecturers in the RAND course, especially in connection with this first section, Techniques of Systems Analysis, because this lecture was prepared to illustrate points they had made earlier. These lectures and lecturers are:

Description of the Course	E. S. Quade (Course Director)
Analysis for Air Force Decisions	C. J. Hitch
Missile Systems for the Future	E. S. Quade
The Why and How of Model Building	R. D. Specht
Behind the Briefing: Inside the Black Box	E. S. Quade
The Relevance of Costs	M. W. Hoag
Objectives	B. H. Klein
RAND and the Air Force	J. R. Goldstein G. H. Putt
General Issues of Criteria	R. N. McKean
Attrition	G. H. Clement

Game Theory	M. Dresher
Some Special Valuation Problems	M. W. Hoag
Dollar Costs for Air Force Systems	G. H. Fisher
Generalized Technical Characteristics	L. B. Rumph
Behind the Models Behind the Briefings	A. J. Wohlstetter
The Treatment of Uncertainty	A. W. Marshall
The Phasing Problem	M. W. Hoag
A Better Treatment of the Example	E. S. Quade
Alternative Methods of Analysis	R. D. Specht
The Positive Side	C. J. Hitch

The above lectures are now being edited for publication by E. S. Quade. The present volume should be regarded as a companion to these published course lectures and care has been taken to avoid overlap.

Special thanks are due to George Margadonna and his staff who prepared all the charts, mostly working against deadlines. The thanks is not only for their work and skill but for their at least overt cheerfulness in the face of the many changes we made.

Finally, we would like to thank the people who read and commented on sections of the early drafts. There are enough of these (people, drafts, and comments) so that we are embarrassed that the report did not turn out better. We would like particularly to thank: L. C. Bohn, W. M. Capron, S. T. Cohen, F. R. Collbohm, L. A. Foster, T. E. Harris, B. W. Haydon, O. Helmer, L. J. Henderson, Jr., H. P. Lynn, Jr., W. G. McMillan, M. P. Peisakoff, M. S. Plesset, J. W. Tukey and P. Weidlinger.

RM-

TECHNIQUES OF SYSTEMS ANALYSIS

CONTENTS

RM-182?-1 16

INTRODUCTION[1]

Since World War II, people with scientific training have played an every increasing role in government, especially in the military establishment, not only as technical advisers and designers of equipment but as policy advisers. In particular, they have often been asked to aid in making managerial-type decisions which used to be the sole concern and prerogative of executives and bureaucrats. This kind of staff activity, while going back, at least sporadically, to antiquity, received such an impetus during World War II that we generally say it started then. It often goes under the name Operations Research. As the name implies, it is mainly the study of operations, particularly, but by no means exclusively, military operations.

Today, in addition to Operations Research, there is a similar but broader activity called Systems Analysis. Systems Analysis bears about the same relation to Operations Research as strategy to tactics. They both look at the same sorts of questions, but Systems Analysis is broader (and therefore less detailed) in both space and time. This book is essentially about the broader activity, though there will necessarily be much discussion of the narrower as well.

---

[1]The first part of the book is long. What is worse, some may find it tedious. We apologize for both the length and the tedium, but we know of no way to give the reader the "feel" of a Systems Analysis, except to go through an example of this sort. It is indeed the purpose of this section to furnish the reader with vicarious experience. A more transparently vicarious experience can be obtained by simply reading Appendix III which summarizes the main ideas.

The problem faced by the Operations Analyst, particularly during a war, is conceptually quite simple. He is usually asked to study a situation in which there is given a definite type, or small range of types, of equipment and in which there is a definite objective to achieve. The enemy also has definite and fairly well-known types of equipment and objectives. The problem is to decide how to use one's own equipment in the best possible way, given this rather precisely stated environment. Typical problems that were treated in the last war are:

the number of ships to put in a convoy

the spacing between these ships

the ratio of escorts vessels to cargo carriers

evasive tactics of planes, ships or submarines

the optimal aiming of aircraft guns, bombs, depth charges, and torpedos

bombing altitudes

search and reconnaissance

radar operating techniques

There has been a great deal written about the success of Operations Analysts in treating these problems during the war. Even allowing for pardonable exaggerations, the analysts often came up with suggestions that were quite different from practice and yet often demonstrably better. This, in spite of the common and well-founded belief that in the past "experience" has been a better guide than "theory" in this kind of work.

The reason for this success is fairly clear. We might, for example contrast the situation during World War II with that during the Napoleonic

Wars, and ask ourselves if scientific personnel could have contributed much to Nelson's conduct of the battle of Trafalgar. The following quotation contrasts the situation Nelson found to that faced by the professional officer today. It also indicates some reasons why a civilian analyst is sometimes in a better "psychological" position than the professional military officer in approaching new long range problems.

> The professional officer, stimulated always by the immediate needs of the service to which he devotes his life, becomes naturally absorbed with advancing its technical efficiency and smooth operation. This task has become ever more exacting with the increasing complexity and rapidity of change of military technology. Nelson, whose flagship on the day of Trafalgar was forty years old yet in no wise inferior in fighting capacity to the majority of the ships engaged, could spend his lifetime learning and perfecting the art of the admiral without fearing that the fundamental conditions of that art would change under his feet. Today the basic conditions of war seem to change almost from month to month. It is therefore difficult for the professional soldier to avoid being preoccupied with means rather than ends, especially since his usefulness to his immediate superior hangs upon his skill and devotion in the performance of his assigned function. And if there is one thing above all that distinguishes the military profession from any other it is that the soldier always has a direct superior.[1]

Nelson and his contemporaries could have had forty years' experience in handling the equipment they were fighting with. In fact, they could have had effectively even more. They could draw not only on their personal experience, but on the experience of others through personal contacts or writings. Under such circumstances the analytical process does not usually yield results which will compete

---

[1] Bernard Brodie, Strategy versus Tactics in a Nuclear Age, Brassey's Annual, 1956. We have put in the underscoring.

with those that can be obtained by an experienced man using ordinary judgment and inventiveness.

World War II was quite different. Nobody really had much wartime or even peacetime experience with the equipment because so much of it was relatively new. In some cases it did not exist before the war and the military didn't even have the benefit of exercises and discussion. Under such circumstances, a theoretical or analytical approach, particularly if it can be made quantitative,[1] will often prove to be fruitful. In fact, it was found that almost any honest, technically-competent person could turn out worthwhile and interesting results.

This last statement does not carry over at all to the much broader problems faced in Systems Analysis. Here we no longer have a definite context with specific equipment. Sometimes we don't even have definite objectives. Rather we are trying to design a system capable of meeting contingencies which will arise five to ten, and sometimes fifteen, years in the future. We must not only design this system, we must also decide under what conditions it will be used and what we shall want to do with it. The recommendations of the Systems Analyst are mainly concerned with "beliefs," research, development, and procurement, and only incidentally with operations.

Under these circumstances, competent, honest people often don't do very well; this does not, of course, mean that we want incompetent or

---

[1]The fact that many current military operations are practically problems in applied physics is itself one of the main reasons for the "success" of Operations Research. There is some more discussion on this point in the Miscellaneous Comments chapter.

dishonest people. It does mean that, in addition to technical competence and honesty, a certain sophistication is necessary.

To make the role of the Systems Analyst clearer it seems worth while to make a distinction among at least three kinds of conclusions--what might be called intuitive judgment, the considered opinion, and the technical or scientific "fact."

The first is essentially based on the individual's experience and background. It is the basis of the day-to-day decisions of executives, businessmen, and in fact almost everyone. While it may be informed, the machinery by which it has been arrived at is not explicitly shown. It is essentially as good or bad as the man who is making it.

The second we have called the considered opinion. It differs from the intuitive judgment in that the logic behind the judgment is made explicit--this usually means that it is quantitative. In the best case it is arrived at by a reasonable and impartial examination of the known facts with due and explicit allowances being made for uncertainties. In the worst case, it may be an extensive and misleading rationalization of a prejudged position. In both cases it usually claims to be "rational." The value of the opinion still depends on who is making it; however, insofar as the machinery is clearly shown, and not hidden by a mass of charts, calculations, and technical verbiage, the audience has some chance of making its own considered opinion from the information presented.[1]

---

[1]It should be clear to the reader that we have taken some liberties with common usage in making these definitions. For example if somebody spent some days in trying to decide some crucial choice problem and after much internal debate and struggle made the remark, "Well I have done a good deal of cogitating and it is now my considered opinion that I should...," we would probably say he has made an intuitive judgment.

To make the contrast between our definitions of intuitive judgment and considered opinion clearer, it is worth mentioning that in previous times there wasn't much room in human affairs for considered opinions. In most situations there were experienced men available to make off-hand decisions, or the pace of events was so slow that people acquired experience almost without trying. Even when people tried to make opinions explicit, the best they could usually do was essentially a simple or complicated listing of the pros and cons with little or no explanation of how to balance the pros and cons quantitatively. In addition there really wasn't much place for any process of arriving at conclusions that tends to take 3 to 12 months and uses "analytic" rather than "practical" processes, except in the fields of criticism, commentary, or reform.[1] The contrary seems to be true today--hence, a major reason for what is called Operations Research and Systems Analysis.

The last kind of opinion is the scientific or technical "fact." While such "facts" are much more subject to controversy than the general public suspects, it is still true that they can usually be clearly separated from the individual and are in some sense "objective." In particular, insofar as the opinions are based on experiments, logic, or calculations, other people will invariably have repeated the steps and come up with the same answers, or the results will not be believed.

---

[1]There will be some more discussion of this point in Part Three of our book. Roughly the problems were on the whole less amenable to analysis and even when they were amenable the people and techniques were not and could not be expected to be available.

Let us give an example of the three kinds of opinions. Suppose a General asks an Operations Analyst how to increase the number of bombs dropped on the enemy in the next 30 days. The analyst starts by gathering together or discovering a lot of technical "facts." These might concern such things as:

the skill of pilots, bombardiers and navigators
the performance of the planes at various altitudes
the performance of the enemy's defenses (radars, fighters, missiles) at various altitudes
the deployment of the enemy's defenses

He puts these things together as best he can and studies various tactics. If he is good he may include inventions of his own in the field of either equipment or tactics in the study. Assume now that when he finishes his study, it is his considered opinion, and furthermore he can (whether correctly or not does not for the moment concern us) demonstrate how he arrived at this opinion, that the General can

1. double (± 20%) the number of bombs on target at the cost of tripling (± 20%) the number of planes shot down,

or

2. triple (± 20%) the number of bombs on target at the cost of increasing the attrition by a factor of 10 (± 30%).

He finds it too difficult to carry the analysis any further and try to decide which of the two tactics is best, so he reports his findings to the General. The General looks over the analysis, believes it, and after a good deal of cogitating and consultation with an internal crystal ball, decides to try the tactic which triples the number of bombs on target. When pressed to explain his decision, he will say something like, "It seems to me so important to get those bombs on target in the next 30 days that I am willing to lose most of my air force doing it."

Since with the information he has he could have also made exactly the opposite remark, we are forced to conclude that he has made an intuitive judgment. Note that we have not said that his judgment is likely to be wrong, or that there is a better way to make the judgment--only that almost none of the reasoning behind it was made explicit.

There is another important thing to notice. If the General had done the analysis himself, in view of the uncertainty of how to trade pilots and planes against bombs on target, he might have been willing to make a non-quantitative analysis of the relationship between the numbers of bombs on target and planes lost.

However, it would do almost no good at all for the Operations Analyst to go through an almost completely qualitative analysis and tell the General, "Well, I've looked at the situation quite carefully and have concluded that

1. we can put a fair number more bombs on target if we are willing to lose many planes, and
2. we can put a great many more bombs on target if we are willing to lose an extremely large number of planes."

In order to transmit information which can be used as a basis for decision he must quantify the words "fair," "many," "large," etc. and indicate reasonably explicitly the uncertainties.

While the Systems Analyst tries to deal in scientific and technical facts, most of the time he does not get any farther than some combination of these with a considered opinion. It is our belief that if he is doing his job properly he should not regularly originate conclusions (nor be under pressure to do so) that are mainly intuitive judgments, although some will disagree with this. Unless conclusions can be supported by explicit (and ordinarily quantitative) analysis they should be labeled conjectures. In general, conjectures are best sold by reasonable men with experience or by enthusiasts. Insofar as the Systems Analyst does this he should not call it Systems Analysis.

Another reason why the analyst should try to restrain himself from indulging too often in recommendations based only on some intuitive ideas is that he should try to preserve a detached professional attitude. If he is always giving intuitive judgments he is much more likely to become embroiled in day-to-day policy fights. He then not only tends to become partisan but, what is worse, he will lose that feeling of relative leisure and detachment which is often essential to good work. It is not that considered opinions are always or even usually correct or non-partisan, but that intuitive judgments, by their nature, cannot be explicitly justified and, therefore, almost automatically increase the possibility that one will take sides as an enthusiast rather than an analyst.

We mentioned before that sophistication is essential for good work in Systems Analyses--at least those with a broad context. To illustrate it, however, one would not only have to talk about the real world, real problems, and real policies, but also about the real mistakes. (The first three, at least, are too difficult to describe here and the last tends to be kept classified.) Instead, we will talk not so much about Systems Analysis itself as about some of the techniques it uses in arriving at considered quantitative opinions. While we shall try to illustrate as clearly as possible what we mean by a good analysis, we will actually be treating a rather simple and almost trivial problem. In fact, we shall be playing with a toy rather than studying a realistic situation. We hope that the toy is educational, but educational or not, it is only a toy. All the numbers are hypothetical and we present very few substantive results. We must insist on this point. We still bear scars from an unpleasant experience caused by our not making this caveat forcibly enough.

The Problem

Specifically we are going to consider the problem of designing, developing, and procuring a strategic air force. Our study will be divided into the four parts indicated by Chart 1.

THE FOUR PARTS OF THE PROBLEM

I DESIGNING THE OFFENSE

II PROBABILISTIC CONSIDERATIONS

III DESIGNING THE DEFENSE

IV THE TWO-SIDED WAR

CHART I

I. Design of the Offense

We start by considering a series of budgets and allocating the money between planes and bombs. We try to choose an allocation that maximizes the amount of damage that the force we buy can inflict on the enemy. In doing this analysis, we shall use the so-called expected-value model. This picture of the world ignores most of the effects of uncertainty. It assumes that what happens on the average actually does happen. The notion is described and criticized in more detail later.

II. Probabilistic Considerations

Next we go beyond the expected-value model and consider the effects of uncertainty. We re-do the analysis of Chapter 1, and see how our results and expectations are changed by the explicit inclusion of the effects of fluctuations. We will find that for some questions and some situations the previous expected-value model is satisfactory. For other questions, it is not.

III. Design of the Defense

In this section, we consider how to defend our offensive force. In line with certain conclusions arrived at previously, we use an expected-value model as in Chapter 1. However, even though we use the same sort of model the study is more complicated. The complication arises partly because there are more components to take into account but, mainly, because our standards are higher. The reasons for the higher standards are pedagogical and do not reflect any intrinsic difference between offense and defense. We have simply postponed some of the more subtle considerations to this section. In particular we will emphasize the distinction between what we call Statistical and Real Uncertainty.

IV. The Two-Sided War

Here we introduce the most complicated question of all, the action and reaction of the enemy. Our point of view will be so symmetrical that we have called this section the Two-Sided War. We find there are serious compromises between what is actually done in practice and what one would like to do--compromises which are much worse than those made in the first three parts. Probably almost all right-minded people agree with our treatment of the problems discussed in Chapters 1, 2, and 3. We will find, however, that the treatment of Chapter 4 is controversial even among the select members of the profession with whom we like to be en rapport.

CHAPTER 1: DESIGNING THE OFFENSE

Models

In discussing a problem of this type, the first thing one usually does is set up a model of the situation. All the military readers are probably familiar with the terms "concept" and "doctrine." Concept is that which is believed; doctrine is that which is taught. To these two terms, we wish to add a third term, "model." A model is that which is analyzed. Broadly speaking, it comprises the assumptions of the study. (It should be made clear that while the exposition may start with a description of the model, the decision as to what model or models to use usually comes late in the study after a good deal of work and insight.)

Very often staff papers will start with a listing of assumptions. Then one may find upon reading the paper that many of the assumptions that actually influenced the results were not listed while many of the listed assumptions were either irrelevant or were ignored.

If a mathematician or scientist makes a model of a situation, he is generally successful in making the assumptions he is going to use very explicit and then in relating his conclusions to these assumptions in a fairly direct way. It does not follow that he is necessarily better in making assumptions--only that his mistakes will be more evident. One can be technically good at deriving conclusions from assumptions and yet very poor at making realistic assumptions. To quote Ovid, it often happens that "the workmanship is better than the materials." We shall illustrate this model-making process in what follows. Later, there is some general discussion about the use and abuse of models.

As shown in Chart 2, there are five important components in our model. We shall assume that we can buy the first two, planes and bombs,

ELEMENTS OF MODEL

1. PLANES
2. BOMBS
3. AREA DEFENSE
4. LOCAL DEFENSE
5. SHELTERED AIRFIELDS

CHART 2

for $5,000,000 apiece. We shall defer the discussion of the costs of the next three items to Chapter 3, Designing the Defense. The Area Defense is that part of the defense which can be deployed tactically to meet threats when and as they appear. Local Defense is that part which can meet threats specific to only a small geographic area. Finally, the sheltered airfields house our targets, the enemy's planes.

In order to illustrate the nature of model making, we shall consider two different models of the Area Defense. The first model will be a fairly realistic one used in many RAND studies. However, because it is realistic, it is too complex for our purposes. We shall, therefore, also consider a much simpler one that we will actually use. By considering two different models we hope to illustrate some of the flexibility possible in model-making.

Area Defense Model I

Consider an area defense model as used in some RAND-type studies. This model is really a map exercise which simulates an actual attack in some detail. Conceptually, what we would like to do is to run off a series of wars, send our bombers over enemy country, drop bombs, let them shoot at us and see what happens. After this exercise, we would restore everything to its original condition and try another experiment with a different system. However, even fanatics for the experimental method agree that this particular operation is unfeasible. Actually, any exercise with real equipment may be unfeasible, as the equipment that we want to test may not even be in the drawing board stage. We are then almost completely restricted to pen and paper exercises.

Let us consider how one might run through such a pen and paper exercise. The steps are illustrated in Charts 3 and 4.[1]

AREA DEFENSE MODEL I

1. DEPLOYMENT	6. COMMITTAL
2. DETECTION	7. ABORTING
3. TRACKING	8. GROSS ERRORS
4. IDENTIFICATION	9. INTERCEPTION
5. AVAILABILITY	10. AIR BATTLE

CHART 3

[1] Chart 4, which is not included here, is an animated version of Chart 3.

First, we deploy our forces; that is, we make up a war plan with definite forces and tactics. Then we deploy the enemy's forces; that is, we decide where to locate his radars, ground observers, airfields, local defenses, missiles, communications, airborne patrols, and all the other things that go into a modern air defense system. Only after both forces have been deployed can we start the calculation of the strike.

The strike consists of a series of events most of which have probabilities associated with them. The first is the probability that the enemy will detect one of our planes--recognize by radar or by ground observers that an unknown object is in the skies. He must then track this object, getting its speed, course, and altitude. While this tracking is going on, he will attempt to identify the unknown as friend or foe. If he cannot identify it, he must have available an interceptor to send up and attempt either to check by visual scrutiny whether it is hostile or, if he is certain that it is hostile, to shoot it down. There must be a committal policy; he must decide in advance under what conditions he will send up planes, and how many planes he will send. There is a definite probability that his interceptor will abort, usually because of an equipment failure. There is a possibility of a gross error. The pilot or controller may, for example, make a complete and utter "butch." Then, the plane must make an interception; that is, even if his plane gets to the correct target area, it may or may not actually make a contact with our plane. Finally, there is the air battle. Who, if anybody, gets shot down?

Area Defense Model II

Now, as one might imagine, running through such a simulated pen and paper map exercise can be a long and tedious process. It may take days, weeks or even months, according to the size of the attack and the detail with which it is simulated. It is, therefore, much too time-consuming for the kind of study we are going to describe. We shall, instead, consider a much simpler but still quite useful model of an area defense system. It is shown in Chart 5. Here, we are simply given a curve that gives the

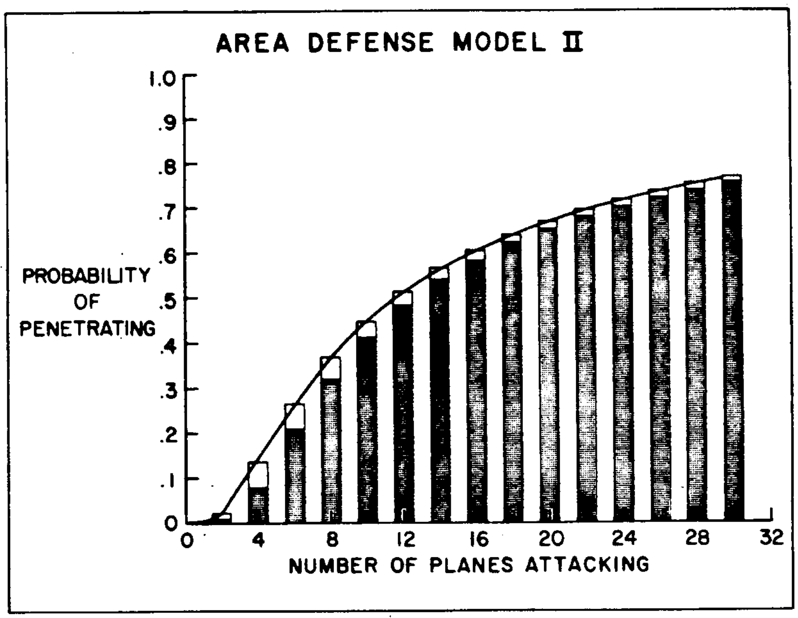

CHART 5

probability that any individual plane will get through the enemy's area defense.[1] This probability is a function of the total number of planes that attack. For example, if we refer to the chart we see that if 8 planes attack, each will have a probability of roughly .36 of getting through. If we send 16 bombers, then the probability of each plane's getting through is roughly .6. In general, the more planes we send the better the chance for any particular plane to get through.

At first sight it may seem curious that adding a plane increases the probability that another plane will survive. This occurs because of the phenomenon known as saturation. Sending more planes confuses and overextends

[1]The mathematical model behind this and other basic data charts is described in Appendix 2.

the enemy. It increases the chance that he will make a mistake or be fully committed. If, for example, the enemy has just two interceptors and we send just 1 plane, he would find it very simple to direct both of his interceptors to that 1 plane. If we send 2 planes, not only is he limited to sending only 1 interceptor against each plane, but because he is trying to keep track of 2 interceptors and 2 planes simultaneously, his efficiency in keeping track of any one of these objects is reduced. His chances for making mistakes are increased and he has to divide his effort.

The unshaded portion at the top of the bars in Chart 5 gives the extra amount of probability of penetration that the last plane we sent bought us. We have redrawn these unshaded portions of the bars on a separate chart so that we can investigate how this increase in the probability of penetration changes as the number of planes changes.

Chart 6 shows that at first each extra plane sent buys more extra probability of penetration (for all the planes) than the previous one did.

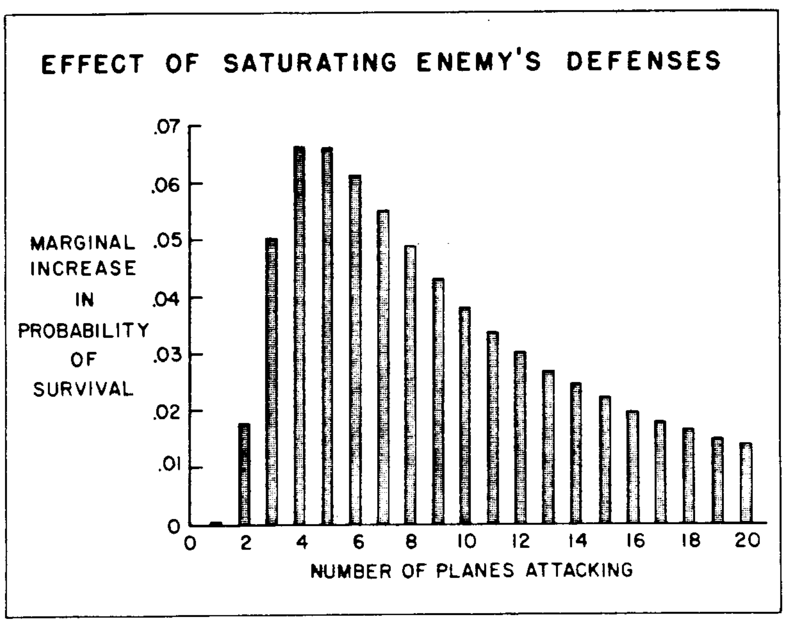

CHART 6

From five planes on, however, each plane buys a smaller increase than the previous one bought. At about 20 planes, the extra probability gained by adding one more plane is only .015.

This situation is familiar to the economist. He says that if we send 5 planes or less, we are in a region of increasing marginal returns; that if we send 6 or more, we are in a region of decreasing marginal returns. That is, the marginal return of investing one more plane, as measured by the probability of getting through, first increases and then decreases. This qualitative behavior will be important to us later.

It should be mentioned that at this stage it is more useful to consider the probability of getting each plane through rather than the expected number of planes for each attack. One reason for this is that when the planes are different (as in the case of bombers and escorts) they must often be treated separately (as in Chapter 2).

It is clear that describing an area defense system by Chart 5 is much simpler than carrying through the detailed map exercise mentioned previously. Whether it is reasonable or not so to represent an area defense system depends on the tactics that are being studied, the kind of analysis that one is doing, and whether one has the right kind of crystal ball from which to get the curve. In particular, a single curve can be expected to approximate only a very limited range of circumstances and one might need very different curves if conditions changed. Let us, however, fix on this single curve and go on with the study.

Local Defense Model

Once we have penetrated the area defenses we must penetrate the local defenses. In our study, we will also characterize the local defenses by a single curve that gives the probability of penetration as a function of the size of the attacking force. This curve, shown in Chart 7, also has a region of increasing and decreasing returns.

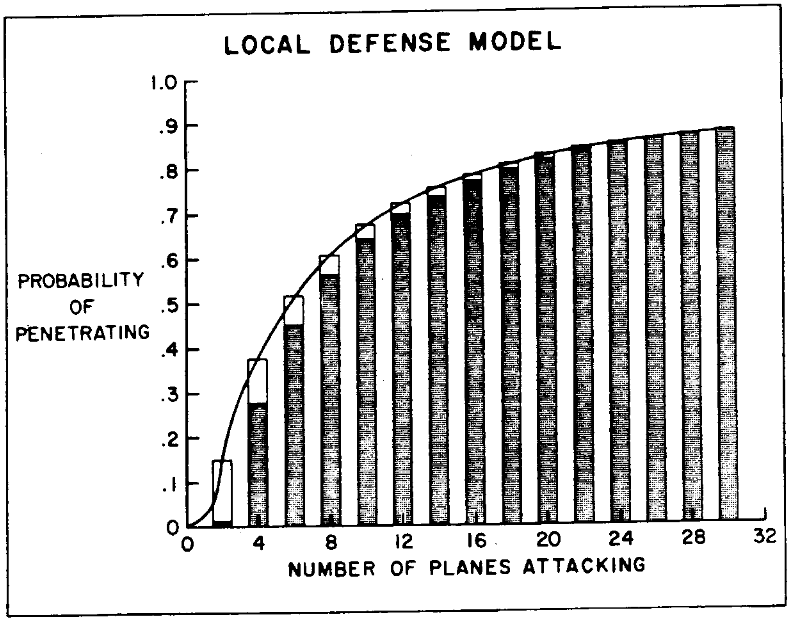

CHART 7

Airfield Model

We now come to the last element of our model, the target airfields. We are going to assume that the enemy has only two. The reason for this assumption is not that we don't know he has more, but that we want to keep the study simple. We need to give the enemy at least two because this is the minimum number which will allow us to demonstrate the effects of dispersion. The enemy's airfields can be completely characterized for the purpose of our study by giving the average percentage of the planes on each field which are destroyed by any given number of bombs dropped on the field. Chart 8 does this. By referring to this chart we see that one bomb will destroy 40% of the airplanes on the field; the next bomb destroys 24% more, making a total of 64%, and so on. In this case, we have decreasing

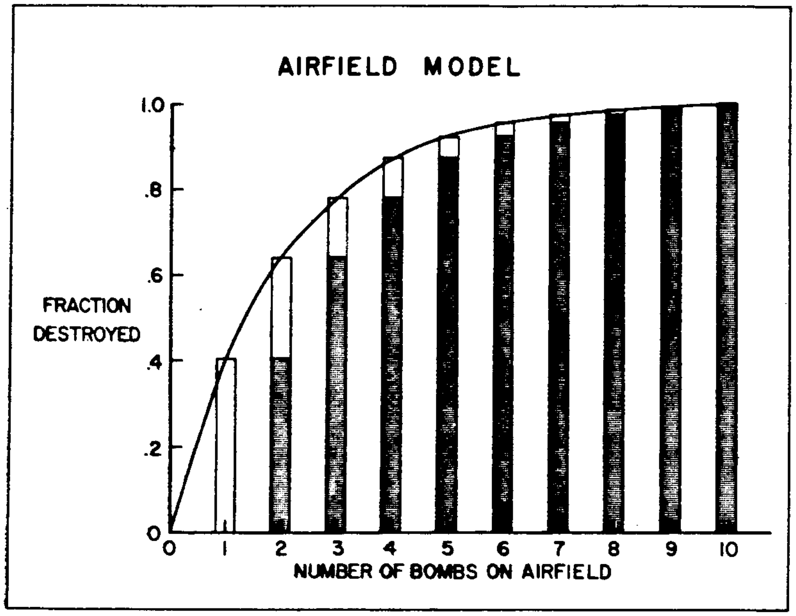

CHART 8

returns everywhere; that is, every bomb destroys--we are always speaking in terms of on-the-average here--less than the previous one did. The reason for this is that once we have dropped a bomb on the field, there is a certain probability that the area destroyed by the next bomb will overlap an already destroyed area. Therefore, this second bomb will, on the average, waste a part of its energy in destroying an already destroyed place.

Charts 5, 7, and 8 can be considered as giving the numerical assumptions of our illustrative study. All of our quantitative conclusions are based on the curves on these charts.

Tactics

Now how shall we use our air force? What will our tactical doctrine be? Actually most of our tactical decisions were settled as soon as we drew our attrition curves. That is, Charts 5, 7, and 8 assume definite operational tactics. About all that is left to decide is how to allocate our forces to each of the enemy's airfields. In principle, we should consider all possible allocations and then pick the best one. This turns out to be too general and difficult an approach, so we shall begin by considering only two. This is not bad if these two ways are representative enough to typify reasonably all the possibilities. We shall check at the end of this section and see if we can improve our capabilities by operating in other ways. By restricting ourselves initially to only a few reasonable methods, we can easily and cheaply learn about the problem and train our intuition.

The two methods of operating the force are illustrated in Chart 9. We call them a Type I and Type II tactic, respectively. In the Type I operation, we send some number of planes, say 14, against the area defense.

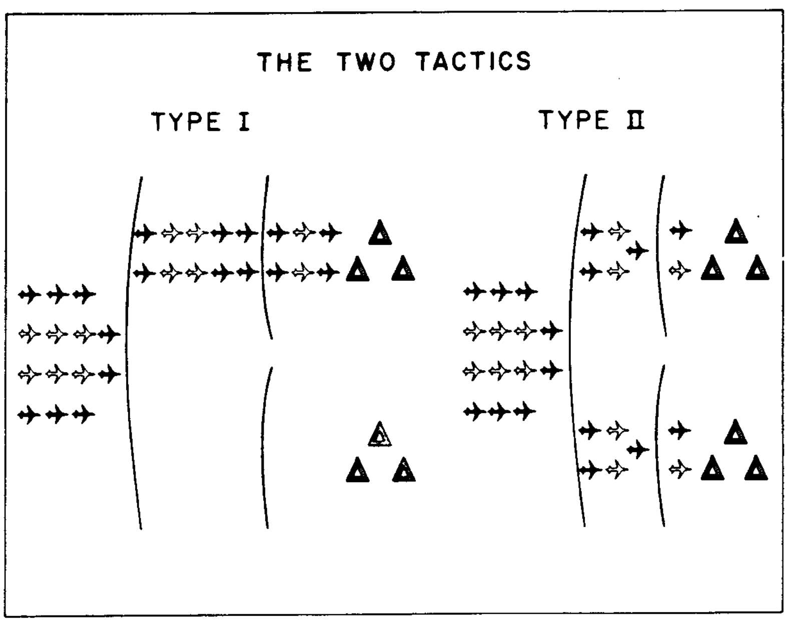

CHART 9

Of these 14 planes, 8 (drawn in dark) are escorts and only 6 (drawn in light) carry bombs. The reason the escorts do not carry bombs is that bombs are (in our model) very expensive, and we prefer to buy extra planes which will confuse and saturate the enemy's defenses rather than to put a bomb in every plane.

Of these 14 planes, we have shown 10 getting through the area defenses, of which 4 are bombers and 6 are escorts. In the Type I tactic, we take these 10 planes and deploy them all against a single airfield. In this case, 4 escorts and 2 bomb carriers get through the local defense. The 2 bomb carriers then drop their bombs on the local airfield. For the present, we shall stop the analysis at this point--we shall not follow them home.

Type II operation starts out the same way--for example, let us say again that 14 planes attack the area defenses and 10 get through. This time, though, we split the surviving force into two equal parts and each part attacks an airfield. Now,we have a smaller force attacking each airfield, so a smaller percentage survives. Notice that it is not only a smaller number, but also (because of the previously described saturation effects) a smaller percentage. If we are lucky, we may get a bomb carrier through each of the local defenses and get to drop one bomb on each airfield. This illustrates the two general methods of operating the air force.

Specializing and Simplifying the Model

Now notice that,as long as we restrict ourselves to two methods of operation, we can summarize the information given in the first two original assumption charts into a single more convenient one. We do this by combining the separate probabilities of getting through the area and local defenses into a single over-all probability of survival for each operation type. The result is shown in Chart 10.

The chart is calculated in the obvious way. Assume, for example, we send 16 bombers against the area defense. The probability of any particular bomber's getting through is roughly .6 (see Chart 5), so on the average

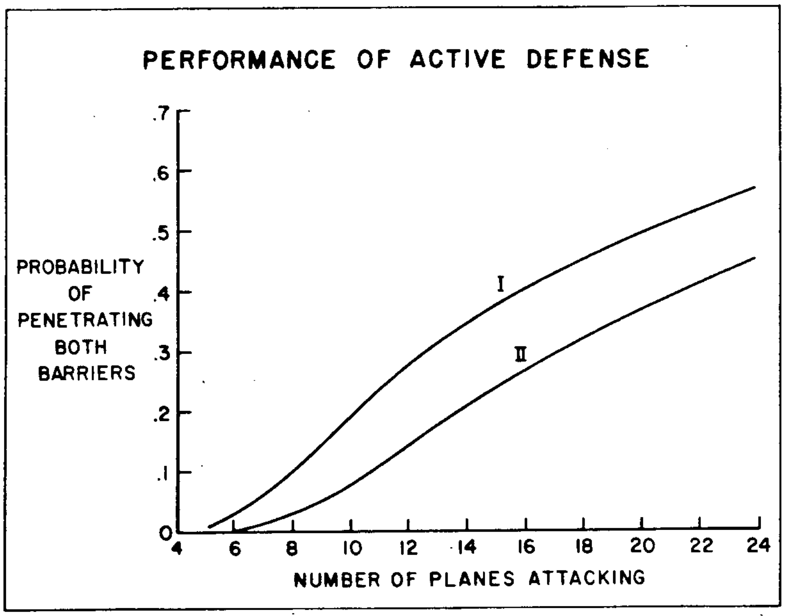

CHART 10

.6 of 16 or 9.6 bombers will get through. Since we are using the expected-value model, we will calculate as if 9.6 bombers actually do get through. We are doing a Type I operation, so all of these 9.6 bombers are sent against one airfield. If we now look at Chart 7, we see that each of these 9.6 bombers has a probability of .68 of getting through the local defense. Multiplying these two probabilities, (.6 and .68) we get an over-all probability of .41 of getting through. This point can be read on Chart 10.

The second curve on the chart is the over-all probability of getting through with a Type II operation. It is computed in a similar fashion except that instead of sending the 9.6 surviving bombers against one field, we split them into two forces of 4.8 bombers each, and send each of these forces against a field. By referring to Chart 7, we notice that each of these 4.8 bombers has a survival probability of .38. Since the original probability of survival through the area defenses was .6, the total penetration probability is .23. By repeating this calculation for all possible initial numbers of planes, we get the curves shown.

We can also present the information on Chart 8 in a more convenient form, as shown on Chart 11. We have drawn the two curves that show the

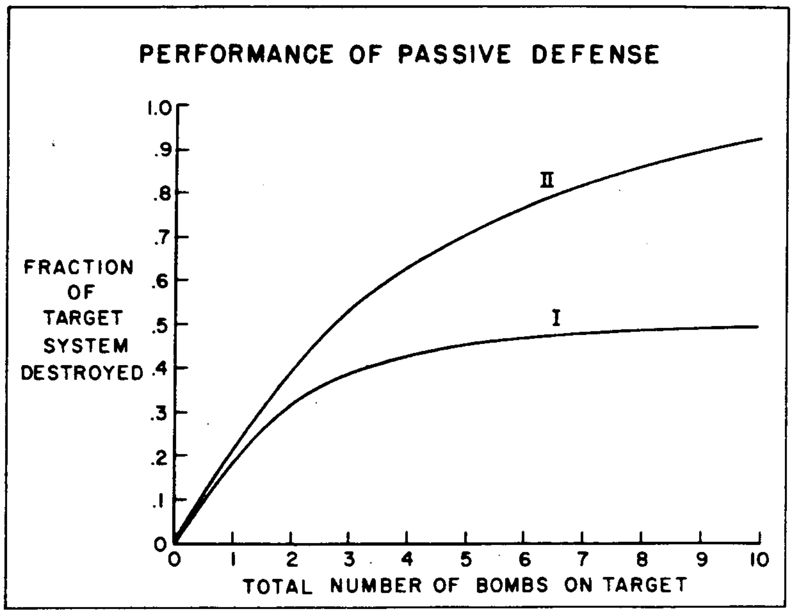

CHART 11

effectiveness of dropping any number of bombs on target in the two types of operations. For the Type I operation, this is essentially a repeat of the curve in Chart 8, divided by two because only half of the target system is at risk on the field that is being attacked. (For example, if we destroy 40% of one airfield we destroy only 20% of the target system.)

In a Type II operation, we simply divide the total number of bombs dropped by two to get the number that actually land on any particular field,

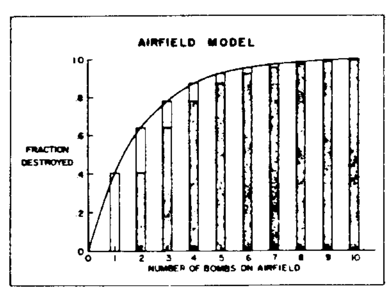

CHART 8

and then find out how much destruction these bombs do. If, for example, 4 bombs were dropped on both airfields, then there would be 2 bombs on each airfield. By referring back to Chart 8 we see that 2 bombs will destroy 64% of a field. Since this is true for both airfields, this is the percentage of the total target system destroyed by 4 bombs.

Because our model is so simple we don't get any spectacular reductions in complexity by specializing in only a few methods of operating. We have, however, succeeded in summarizing the information (Charts 5 and 7 are now on one chart) and putting some of the other information in a more convenient form. In a more realistic study the savings would be more impressive.

If we look again at Charts 10 and 11, we can see very clearly the

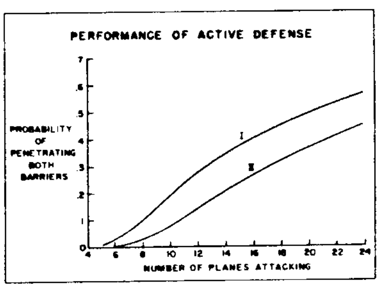

CHART 10

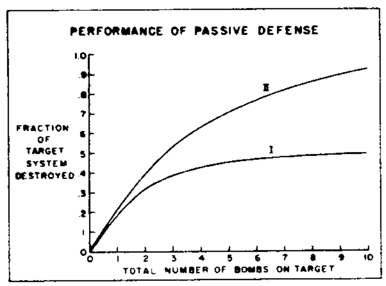

CHART 11

interplay that we shall be studying. The probability of survival in a Type I attack is always greater than in a Type II attack because we have concentrated the attack. For the same reason though, the bombs that we manage to get through are less effective than an equal number delivered by a Type II attack. The question of which is more important--to increase the probability of delivering the bombs or to get more effective use of the bombs which are actually dropped--will turn out to depend on how much money we have to spend and on how we allocate this money between planes and bombs.

Expected-Value Models

We ought to digress and consider the notion of an expected-value model.

In the real world random events occur, and it is usually impossible to say precisely what is going to happen in any particular set of circumstances. However, it is often possible to calculate what happens on the average. The expected-value model assumes (in the calculation at least) that the average result actually occurs.

For example, if we were going to match dimes with the reader in a fair way, we could say that on the average we should neither lose nor win. In the expected-value model one would, therefore, assume that after the toss his fortune was unchanged. As a matter of fact, he would actually be either a dime richer or a dime poorer. However, if he plans to toss many times and his fortune is reasonably large, he will find that it is pretty reasonable to base his behavior on the assumption that he will neither lose nor gain--that is, that the result of each toss is a tie. The alternative would be to calculate in detail the probabilities of all the possible outcomes and then decide whether to play.

To take an example in which the expected-value model won't work, consider the analysis of dice. One can toss a pair of dice and get any result from two to twelve. However, even though it turns out that the average of the results of a large number of throws is close to seven, it would be financial suicide to go into a gambling casino and assume that this average result will occur every time. One would soon learn (rather painfully) all about fluctuation phenomena. The only way to analyze this game is to calculate the probabilities of each separate type of event and evaluate the game in the light of these probabilities.

One of the reasons an expected-value model works out less satisfactorily (in our analysis) for dice than for coins is that when the toss is close to seven (for example a six or eight) the player's winnings may be quite different from what they would be if a seven had come up. We shall find out in Chapter 2 that even (as in the matching game) when the expected-value model computes the average result reasonably accurately it can still be misleading, particularly if one considers realistic objectives. Let us for the present, however, not worry about the errors introduced by the use of expected-value models and continue with our expected-value calculation. (What follows is an example that was worked out in class by the students.)

Numerical Calculation of an Allocation

Assume that we have a fixed sum of money, say $100,000,000, to spend. With this $100,000,000, we can buy twenty objects, each costing $5,000,000 apiece. For example, we could buy 10 planes and 10 bombs or we could buy 12 planes and 8 bombs and so on, as shown by the first two columns of Chart 12. Now, for each allocation of our resources between planes and

WORK SHEET FOR FIXED OFFENSIVE BUDGET

PLANES	BOMBS	TYPE I ATTACK			TYPE II ATTACK		
		PROB. (P)	BOMBS ON TAR. (pB)	FRAC. DEST.	PROB. (P)	BOMBS ON TAR. (pB)	FRAC. DEST.
10	10	.18	1.8	.30	.08	.8	.18
12	8	.27	2.2	.34	.14	1.1	.24
14	6	.34	2.0	.32	.21	1.3	.28
16	4	.40	1.6	.28	.27	1.1	.24
18	2	.45	.9	.18	.32	.6	.13

CHART 12

bombs, we must choose a method of operating the force, Type I or Type II. This is done by calculating the amount of destruction that can be done by each force with each type of operation. The results of these calculations are shown in the last six columns of Chart 12. (The reader should verify them if he is unfamiliar with this kind of calculation.)

The results of the calculations are graphed on Chart 13. As we can see, no matter how we allocate money between planes and bombs, it is always

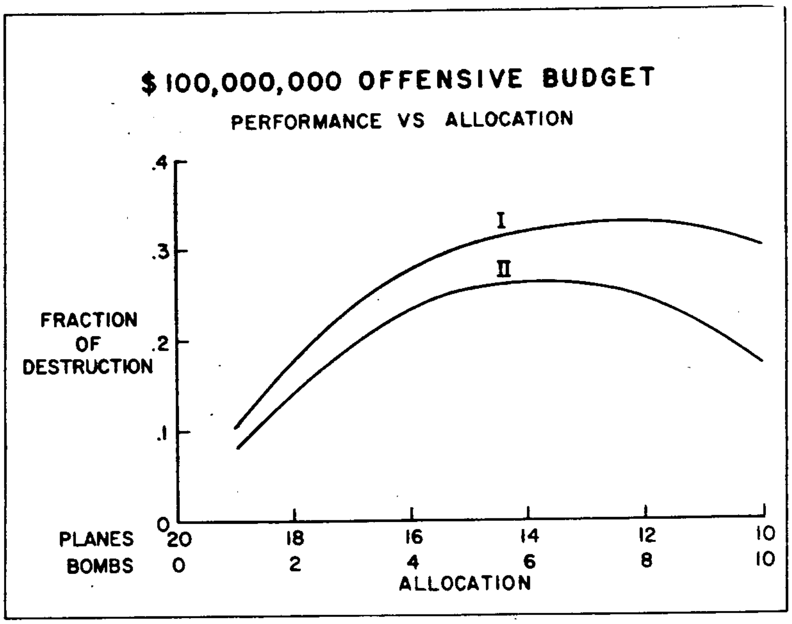

CHART 13

best in this case to do a Type I operation; that is, to operate tactically so as to increase survival rather than to maximize the effectiveness of the bombs on target. The qualitative reason for this rule is that with $100,000,000 we are relatively poor and cannot afford to dissipate our efforts by trying to attack two targets. We also notice that the optimal force composition is about 12 planes and 8 bombs.

We should also notice, and this point is very important, that the curves are fairly flat. If we had chosen a force composition anywhere between

14 planes-6 bombs and 10 planes-10 bombs, we would have lost little in effectiveness.

Also notice, now, an unexpected fact. Suppose we had mistakenly assumed that a Type II operation was better and optimized on this basis. If we later find out that a Type I operation is better, our force is still just about as effective as one which was designed on a correct basis. In other words, while it is fairly important to operate properly, our performance is, at least in this case, not sensitive to the allocation. It isn't even necessary in making the allocation to decide the proper operation of the force. As we shall see in the next paragraph, this insensitivity to allocation is not completely accidental.

Insensitivity by Design

It is worthwhile at this point to digress slightly to investigate in some numerical detail the effect of changing the performance of the planes and bombs. In particular we would like to illustrate how the shapes of the curves that we are studying are affected by the equipment one uses. Assume, for the sake of simplicity, that the relationship between cost and performance, for both planes and bombs, is one of direct proportion. This means that we buy the same performance for the same money--only the package is changed. Thus, if we double the cost of the bomber, we double its performance; it acts like two planes in its effect on the enemy's defenses. Similarly, if we double the cost of the bomb, we double its destructive power. The results of doubling and halving the cost of the bombs and correspondingly halving and doubling the costs of the planes are shown on Chart 14. Curve A shows what happens if we halve the cost of planes and double the cost of the bombs. Curve B is simply copied from Chart 13.

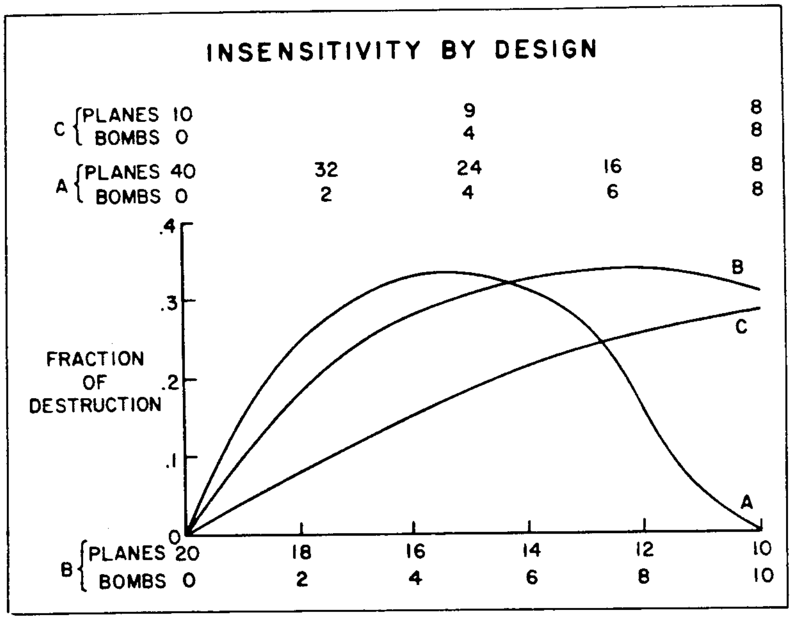

CHART 14

It shows the performance with the original equipment. Curve C shows what happens if we halve the cost of the bombs and double the cost of the planes.

The maximum value that each curve attains is not really affected by these changes, but the shape is. A is narrow. If we were operating with this curve, we would find that our performance is relatively sensitive to the exact allocation of the budget between bombs and bombers. C does not really obtain its maximum because it runs into a boundary set by the fact that each plane can carry only one bomb. Both curves A and C are, therefore, undesirable--in fact, as we shall see, very undesirable. Therefore, if we can choose freely among these curves, we would pick Curve B, and this is what we have done.

In other words, <u>it is no accident that the curves are flat</u>; we go to a great deal of trouble to pick our equipment so that the performance will not be sensitive to the allocation. The question of exactly how to do this will turn out to be crucial. We shall spend quite a bit of time discussing it in Chapter 3. We shall find there that it may even be profitable to accept a lower maximum if this is the only way to make the curve flat.

Back to the Allocation Calculation

Let us consider a slightly bigger budget, $150,000,000. The results are calculated in the same way and are shown in Chart 15. We now notice that a Type II operation is best. We are rich enough so that we don't have

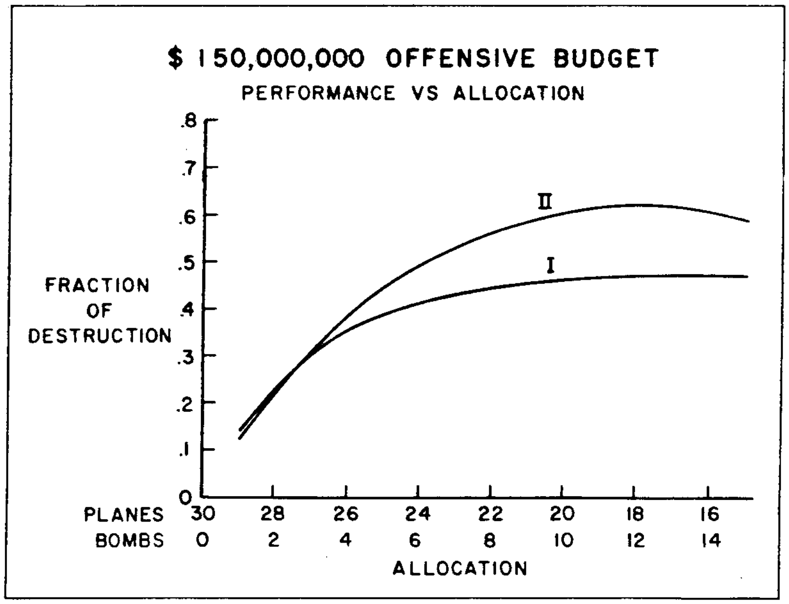

CHART 15

to concentrate our efforts. We can afford to attack the entire target system. Once again, we notice that the curves are flat; exactly how we devide the money between planes and bombs is not too crucial. However, it is fairly important to choose the proper method of operating the forces.

Let us now consider what happens with a really big budget, for example, $250,000,000 (Chart 16). The curves are still fairly flat. We notice, however, another interesting effect. The upper curve has not yet started to turn over. This behavior implies that if you could you would send more bombs than planes; to put it differently, you should be willing to buy lower

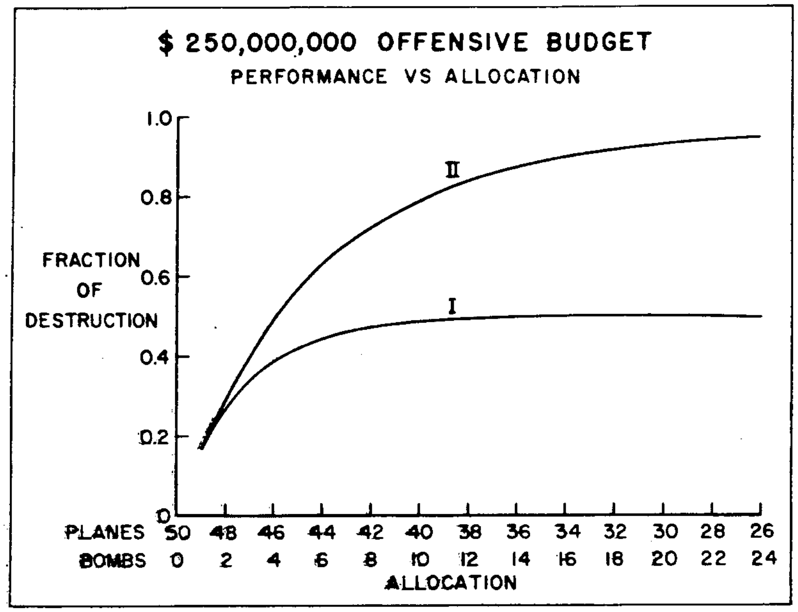

CHART 16

performance (cheaper) planes and higher performance (more expensive) bombs. The situation is that we have sent so many planes that we simply overwhelmed the enemy's defense--in the model. Because the enemy is saturated, these planes are in effect being used for a trucking operation; the enemy's defenses are being circumvented by mass rather than by high performance. We have the slightly paradoxical result that because we are rich we can afford to buy poorer planes. In actual fact there are many situations where the highest performance equipment is not necessarily the best from the system point of view.[1] There is usually no point in using a Cadillac as a pickup

[1]It is probably worthwhile to mention an important psychological difference between the operational people and the general staff that occasionally gives rise to misunderstandings. If you ask the former "What kind of plane should the Air Force buy?", they tend to answer as if the question had been, "Other things being equal, what kind of a plane would you personally like to be in if you were flying over enemy territory?" The answer then is usually, "The highest, fastest, and longest range." The planning staff, which had to look at the overall problem, is sharply aware that other things are not equal and may come up with a different choice. Since, however, the planning staff people are recruited from the ranks of the operational people they have some tendency to have this kind of operational bias.

truck. (In practice it depends a little on what commodity is being "picked up.")

Optimal Allocation and Performance

We have done the allocation problem for all possible budgets. The results are summarized in Charts 17 and 18. The first gives the optimal

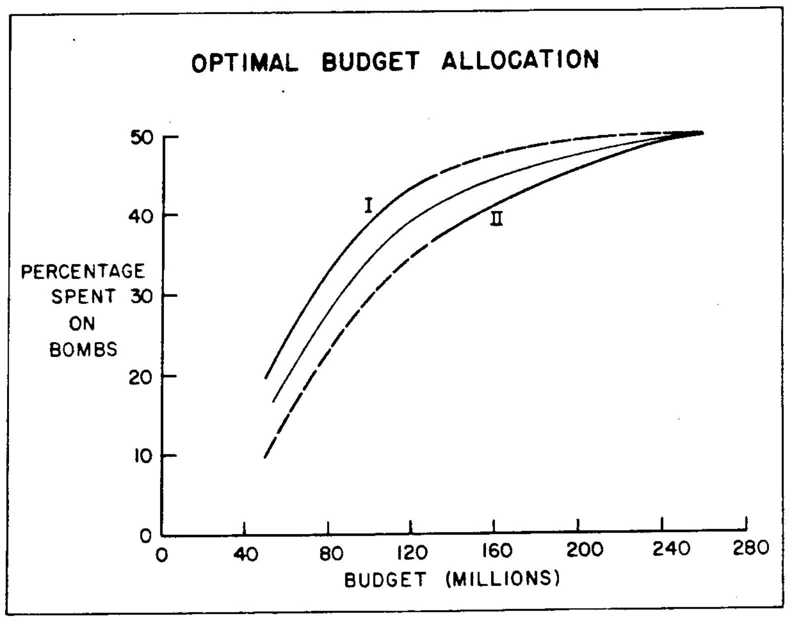

CHART 17

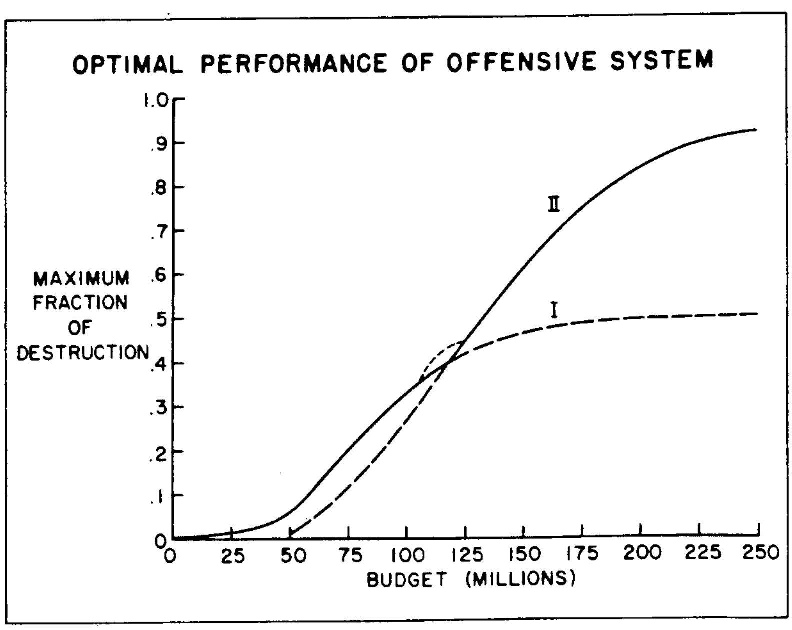

CHART 18

allocation of our money for any particular budget and the second indicates what performance we can expect if we follow this allocation. One immediately notices that the optimal allocation is different for the two types of operations. As a result there is a seemingly sharp change in the optimal allocation when our fortune goes over $120,000,000 and we switch tactics. For example, if we have $120,000,000, we should put about 45% of our money into bombs while, if we have $125,000,000, we should put only 35%.

Presumably we are not fooled by this seeming sensitivity. We should remember that the previous curves were so flat that the allocation could be varied a great deal without changing the results appreciably. Indeed, if instead of using the optimal allocation we use the allocation corresponding to the thin black line between the two curves on Chart 17 (but still use the proper operation type) our performance would drop less than the width of the line given on Chart 18.

It is clear, however, that choosing the correct method of operating the force is important; the differences between the Type I and the Type II operations are relatively large, at least percentage-wise, for most budgets.

Charts 17 and 18 are typical of the end product of most Systems Analyses. They tell you, <u>given certain assumptions</u>, how to spend your money and the kind of performance you can expect if you spend it that way. This is practically the classical Systems Analysis.

In fact, while this kind of study can be very misleading there is a lot of nourishment in it. A sparrow could live on it for a long time. The trouble is that even though there is a lot of information in the study it is not the right kind for policy decisions. <u>Studies aimed at influencing policy should have a completely different flavor about them.</u>

The Central Problem

The difficulty with our study to date is suggested by our title, "Military Planning in an Uncertain World." It is clear that we have done the Planning all right, but it ignores the really hard and often most crucial problem: How does one take account of the whole range of enemy tactics? In addition to the unknown acts of the enemy, how does one hedge against the other uncertainties?

Before making some remarks about this question, it is worth considering the problem of budget allocation and operating tactics in a historical perspective. It is probably clear to the reader that any reasonable person, including for example the ancient Greeks, could have followed our qualitative reasoning and understood, that when one is poor

1) most of the money should be spent on decreasing attrition (buying planes)
2) that one should concentrate on one target,

and conversely that when one is rich

1) one should spend more money on bombs because the enemy's defenses are automatically saturated by the number of planes in the attack
2) that one can now afford to attack both targets.

The exciting thing that we have done is to make the above qualitative remarks numerical; that is, we have changed what we called an "intuitive judgment" into what we called a "considered opinion." How exciting this is can be seen from the fact that the ability to make this type of calculation and end up with Charts 17 and 18 is as much of an intellectual invention as the steam engine or the telegraph is a technical invention. In fact, the concepts needed for this kind of analysis were invented in roughly the same time period as these two gadgets were.[1] Moreover, they were not used for this kind of a question until late in the nineteenth century. In fact, it is only in the post World War II period, which saw a great expansion in the intellectual tools, computing ability, and suitable problems for this kind of analysis, that it really became popular as an aid to the military planner. It is to be feared that it may have become too popular. Many people got so excited about the possibilities that they went overboard and claimed entirely too much for the technique.

One trouble was that people did not generally realize that even modern computing methods are not really powerful enough to evaluate complicated systems without the aid of a good deal of skillful "intuitive" supervision and guidance and, even more to the point, that the problems of uncertainty can swamp or negate a good deal of straightforward analysis. In many cases it was necessary to idealize the problem so much to make it tractable to analysis that the resulting considered opinion was less valuable than almost any reasonable intuitive judgment which was based on an examination of the unidealized problem. However, today Systems Analysts are getting to be both more modest about their claims and better at their work. If the trend continues, we may well come out with a match between claims and product.

Taking up the question of uncertainty first, we will somewhat arbitrarily classify uncertainty problems into three categories. In the first

[1]We are referring mainly to elementary classical economics--in particular marginal utility analysis--and also to such mathematical developments as calculus, analytic geometry, concept of function, etc.

category, which we shall call Statistical Uncertainty, we consider events whose probability of occurrence is more or less objective (such as games of chance). In the second category, which we call Real Uncertainty, we consider events to which individuals may attach subjective probabilities but to which there is essentially no hope, at present, of getting much general agreement as to what these probabilities are. The probability over time of a technological breakthrough is a typical event involving Real Uncertainty. Finally there are the uncertainties due to enemy action and reaction. We treat Statistical Uncertainty in Chapter 2, Real Uncertainty in Chapter 3, and Enemy Reaction in Chapter 4. We shall also (mostly in Chapter 4) try to give the reader a feel for how "intuition" is put into a large calculation to make it manageable.

Excursion Into Elementary Economics

Before we leave Chapter 1, we have a certain obligation to discharge. We have already remarked that the difference between Type I and Type II operations tends to be large. This immediately suggests that we have been a little narrow-minded in restricting ourselves to only two tactics. Perhaps if we had been a little more general and, for example, considered an operation in which we compromised between Type I and Type II by attacking the two airfields unequally, we might be able to improve our performance.

As a matter of actual fact, this can happen. But typically, even if it did, its only effect would be to increase our performance by a small amount. For example, the small dotted curve shown on Chart 18 would be typical of the improvement we might expect. However, with the particular assumptions we have made, it turns out that a Type I or Type II operation is always optimal; no other tactics need be considered.

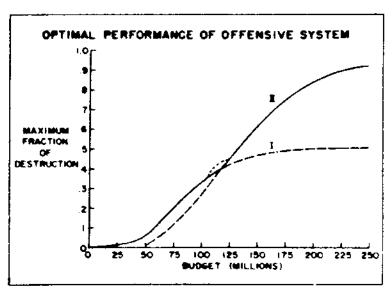

CHART 18

In order to see why this is so, let us consider some of the qualitative characteristics of our curves. It is a little digressive to pursue this somewhat technical consideration, but worthwhile because it will illustrate a type of argument characteristic of this kind of analysis.

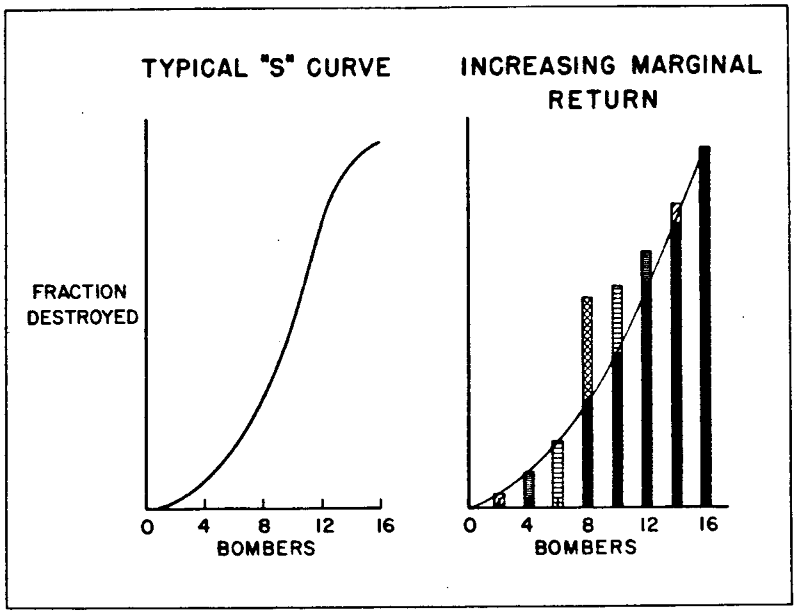

CHART 19

In Chart 19, we have a curve which indicates the value to us (as measured by the average fraction of destruction) of getting different numbers of bombers through the enemy's area defenses. This curve, which takes into account the combined effect of the local defenses and the sheltered airfields, has an "S" shape. The lower part of the "S" corresponds to a region of increasing marginal returns and the upper portion to a region of decreasing returns.

The same chart shows what would have happened if, instead of being "S" shaped, the curve had been everywhere convex downward (had increasing marginal returns everywhere). Assume, for example, that 16 bombers get through the area defense. If we allocate all of these 16 bombers to one airfield, we destroy the fraction indicated. If, however, we take away 4 bombers from the first airfield and allocate them to the second airfield, then we destroy less of the first airfield and more of the second The total destruction is shown as the sum of a dark and a light bar where the abscissa is 12. We have had a net loss in effectiveness. In fact, it is clear that for a curve of this shape, we shall suffer an over-all loss if we take any number of planes away from the first airfield to attack the second one. Therefore, for this type of curve, we should always use a Type I tactic.

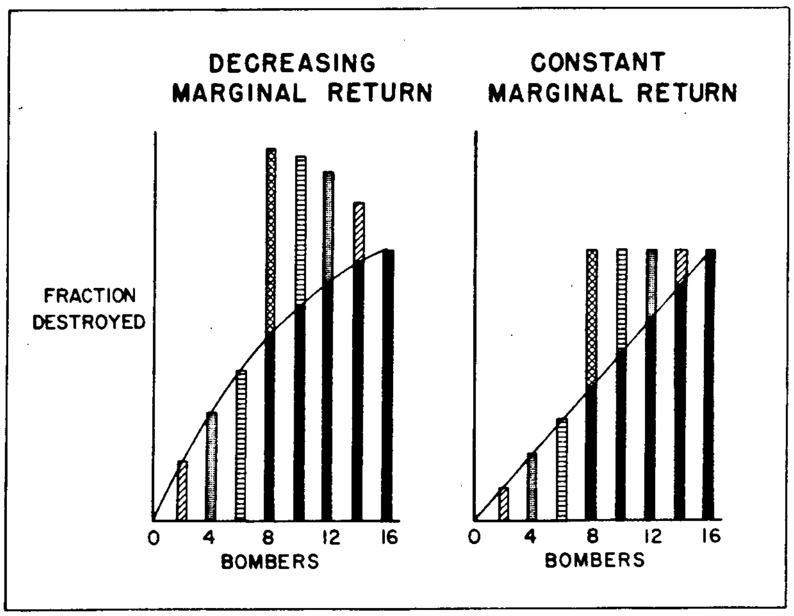

CHART 20

In Chart 20, we show the opposite situation with the curve everywhere concave (decreasing marginal returns). Under these circumstances, we always gain by taking planes away from the first field and attacking the second field until the two fields are attacked equally. Therefore, here one should always use a Type II tactic. For the sake of completeness, the second part of the chart shows that when the curve is a straight line (constant marginal returns), we are indifferent as to how the planes are allocated between the two fields.

Now our problem is a little tricky because our curve is "S" shaped, that is, it has both a convex and a concave portion. In order to understand the effects of this "S" shaped curve, let us look at Charts 21 and 22. In these are presented two typical curves schematized so that we can analyze them. The first schematized "S" shaped curve shows a situation in which the average final slope is greater than the average initial slope. If we are operating at the point of 9 bombers then it does not pay to take away a few bombers from target one and assign them to target two. We shall simply lose in the total fraction of the target system destroyed. However, if we are operating farther out on the curve, as shown in figure (b) of the chart, then it will pay us to trade planes from one target to the other

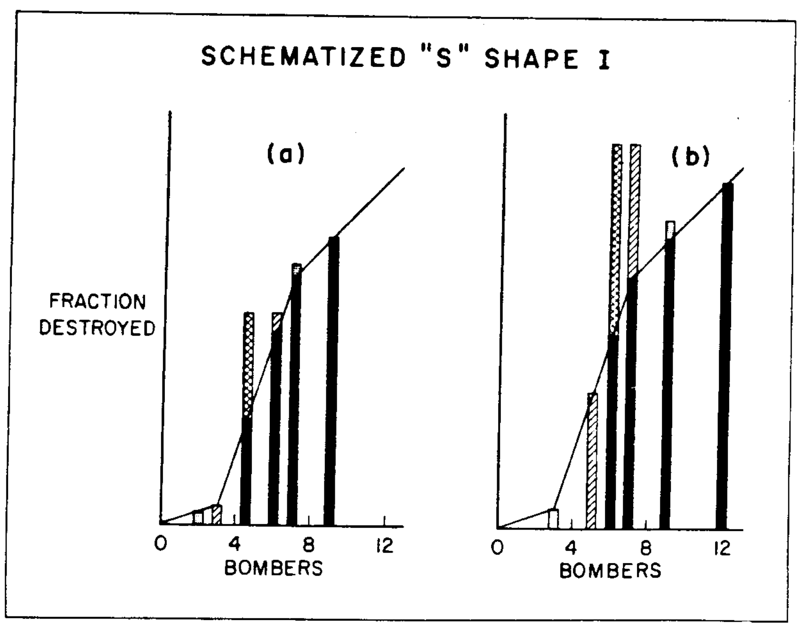

CHART 21

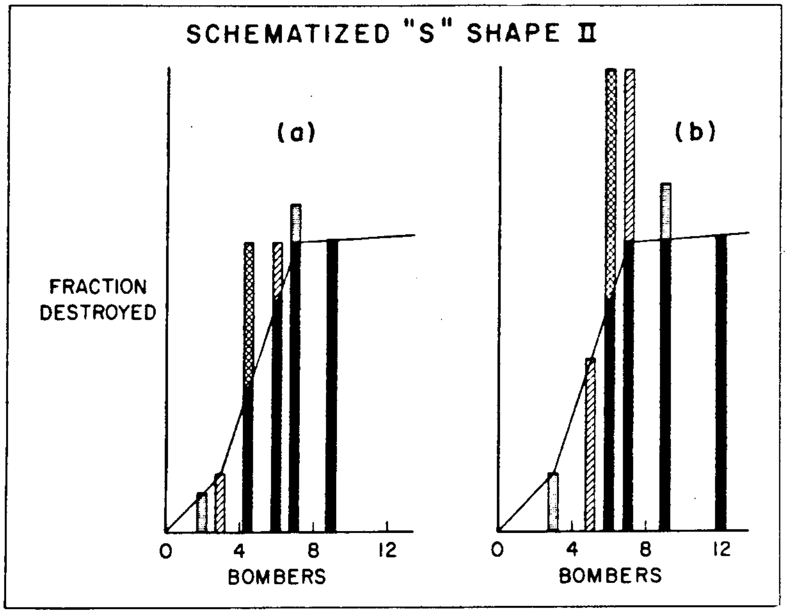

CHART 22

because we then get into the steep middle portion; but we notice that in this case, a Type II operation does as well as anything else. Hence, in this case, either the Type I or Type II tactics are in fact optimal.

Now on Chart 22, let us investigate the case where the average initial slope is greater than the average final slope. The situation is then different. If we are operating at the point of 9 bombers as in the case of figure (a) the 7 bombers on one target and 2 on the other will give us greater destruction than any other combination. On the other hand, if we are operating far out on the curve, Type II is clearly the best. For this kind of a curve and some force strengths, we might want to operate somewhere between Type I and Type II. It does turn out, however, that the "S" shaped curve we are dealing with in our study is not like Chart 22 but is like Chart 21 so that we always want to use either Type I or Type II operation.

CHAPTER 2: PROBABILISTIC CONSIDERATIONS

Why Study?

We are now going to examine the effects of fluctuation phenomena. There are two reasons for doing this. First, an expected-value model actually miscalculates; it is an approximation. Sometimes it is a good approximation and sometimes not. We shall study its accuracy for our model in the first half of this section.

The second, and often more important, reason for using probabilistic (fluctuation) models is that we feel quite differently about things which fluctuate as opposed to things which do not. For example, one would probably react quite differently to the prospect of being given $100 cash or a lottery ticket with one chance in a thousand of $100,000. In the expected-value model these two gifts are identical. In the second half of this section we shall consider how our objectives are, or should be, modified by the explicit introduction of probabilistic considerations.

Qualitative Discussion of Accuracy

It is easy to see qualitatively why expected-value models miscalculate. In Chart 23, we show two curves that purport to give the value to the attacker of getting a certain number of bombers through the area defenses. One of the curves has increasing marginal returns everywhere and the other one has decreasing marginal returns. Consider now a situation in which one has a 50% chance of getting 9 bombers through the area defenses and a 50% chance of getting 3 bombers through. On the average, then, 6 bombers

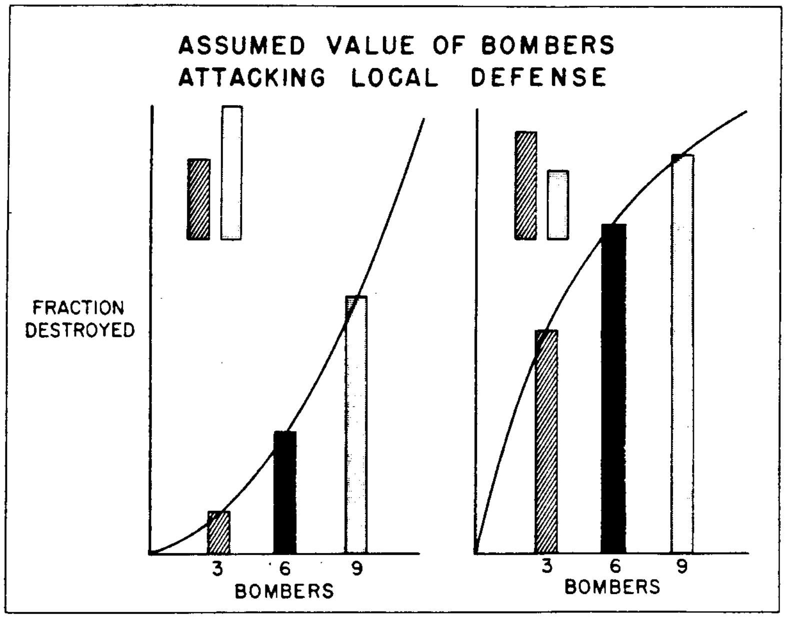

CHART 23

will penetrate. If we are operating with the first curve, the one with increasing marginal returns, we can easily verify that we gain more when we are lucky and 9 bombers get through than we lose when we are unlucky and only 3 bombers get through. In this case, if we use an expected-value model and assume that 6 bombers actually get through, we will underestimate the average damage we do. The amount by which the expected-value model underestimates depends upon how much the number of penetrating bombers fluctuates and the curvature of the curve (rate of increase of the marginal return).[1]

---

[1]Aside to more mathematical readers:

The problem is to estimate the error when the approximation $g(\overline{x})$ is used for $\overline{g(x)}$ where

$$\overline{x} = \int x f(x) dx$$

$$\overline{g(x)} = \int g(x) f(x) dx$$

and $f(x)$ is a probability density function.
Expanding $g(x)$ in a Taylor series around the point $x = \overline{x}$,

$$g(x) = g(\overline{x}) + g'(\overline{x})(x-\overline{x}) + \frac{1}{2} g''(\overline{x})(x-\overline{x})^2 + \cdots$$

(footnote continued on next page)

The second curve illustrates the converse situation where we gain less when we are lucky than we lose when we are unlucky. Finally referring to Chart 24, we can see that the expected-value model calculates correctly situations where the curve is straight (constant marginal returns).

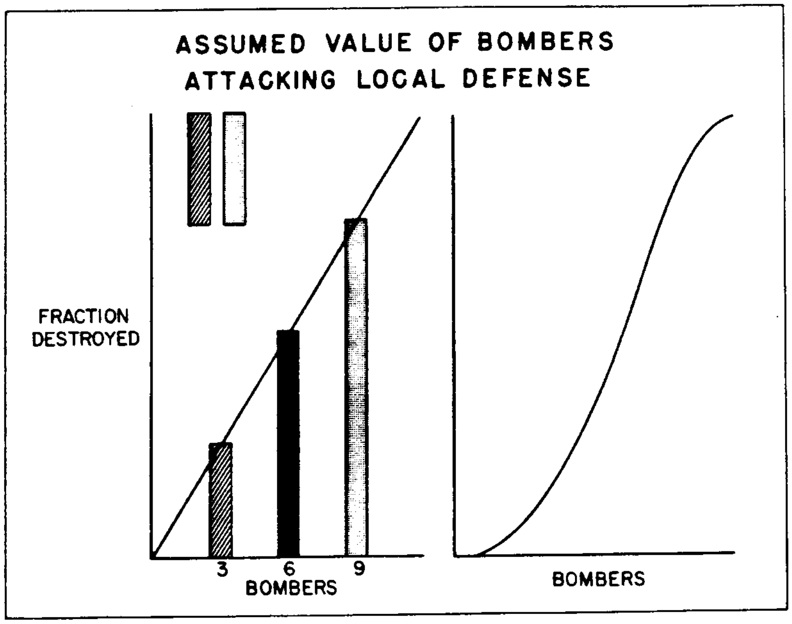

CHART 24

Here the effect of good or bad luck averages out exactly.

We observe that in our problem we actually have an "S" shaped curve; therefore, the expected-value model will underestimate when we are poor and overestimate when we are rich.

---

and substituting into the integral one gets

$$\overline{g(x)} \backsimeq g(\bar{x}) + \frac{1}{2} g''(\bar{x})\ \sigma^2$$

where $$\sigma^2 = \int (x-\bar{x})^2 f(x)dx$$

Since $g''(\bar{x})$ measures the curvature of the curve $g(x)$ at $\bar{x}$ and σ^2 measures the amount of fluctuation associated with $g(x)$, the above formula makes the remark quantitative.

Probabilistic Model

Let us now look at the details of the probabilistic mechanism. Assume, for example, that we have 12 bombers attacking the area defenses. We can see from Chart 5 that each bomber has a probability of .51 of penetrating the area defenses. This means that, on the average, .51 x 12 or about 6 bombers will penetrate. When we used an expected-value model, we assumed that this number actually did get through. As a matter of fact, any number from 0 to 12 might penetrate, depending upon how lucky or unlucky we happened to be.

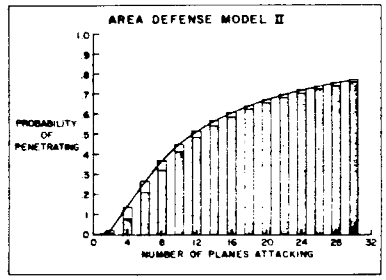

CHART 5

Though the number of planes that get through the area defense is uncertain, the uncertainty is of a kind about which we know a good deal. Once we have assumed that the probability of any particular plane getting through is .51, only so-called Statistical Uncertainty is left. We can,

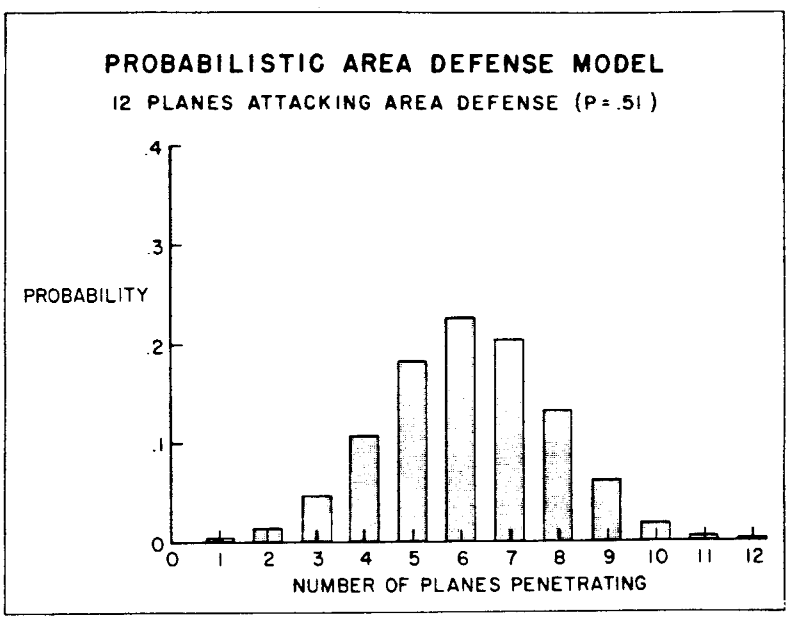

CHART 25

by simple methods,[1] calculate the relative proportion of times (probability) that any particular number will survive. These proportions are shown in Chart 25. While 6 turns out to be the most probable number, still only about 24% of the time does this exact number penetrate and 76% of the time a different number. It is not surprising, therefore, that results calculated on the basis that 6 actually do get through might be misleading.

Chart 25, however, does not give us sufficient information. Assume, for example, that our 12 planes were composed of 8 bombers and 4 escorts. Then, we are interested not only in how many planes get through but also in what kind of planes, that is, escorts or bombers. This kind of information is given in Charts 26 and 27. The first gives the relative proportion

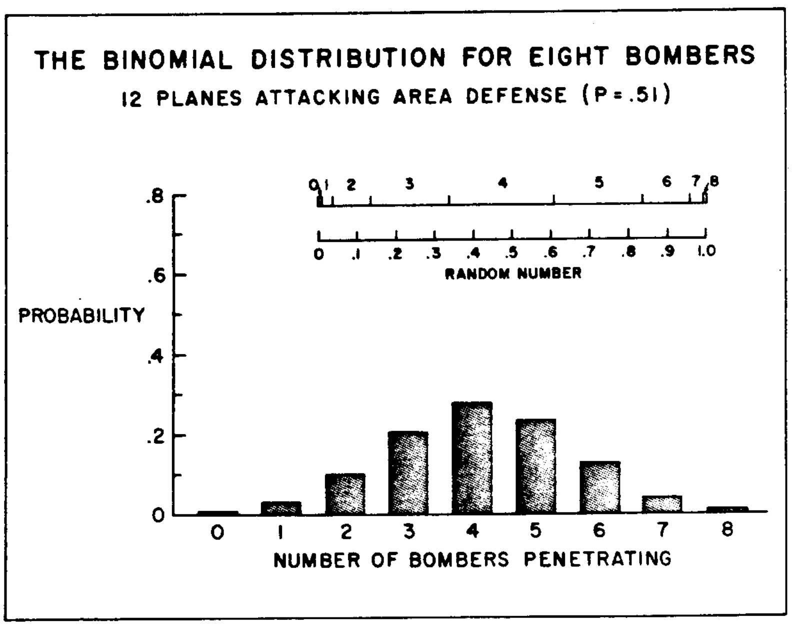

CHART 26

[1]See Chapters 6 and 8 in Part Two on Techniques of Operations Research.

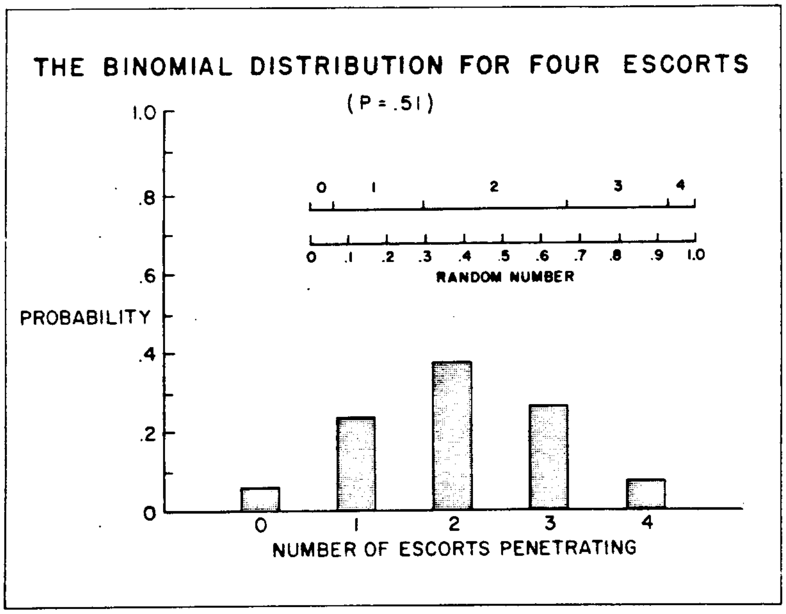

CHART 27

of the time that any particular number of bombers gets through. The latter **gives the relative proportion of the time that any particular number of escorts gets through.**

The Monte Carlo Technique

Taking into account probabilistic effects is simple in principle. The actual computations can, however, be laborious and complicated. In fact, for all but the simplest problem the calculation is too long even for a high-speed machine to do in any economically reasonable, or even unreasonable, amount of time. The truth is that with the simple model we are considering here, we are able, but only barely able, to make the exact calculation. In most problems, we have to resort to a special sort of approximate method of computation known as the Monte Carlo method. The method can be exemplified as follows.

Suppose we have a simple game of solitaire, and are interested in calculating the probability of winning the game. There are two ways to do this.

The first is by exact combinatorial methods. We consider every possible shuffle as being equally likely, and simply count the number of different shuffles that will end in success. If we then divide this number of shuffles by the total number of possible shuffles, we have calculated the probability of winning. A solitaire game does not have to be very complicated before such a procedure would take more than a lifetime, even with high-speed computing machines.

But there is a way to sidestep the combinatorial computation. It is not only easier, but it may be more pleasant also. We can simply play the game, say a thousand times. The relative number of wins then gives us an estimate of the probability of winning. The estimate, of course, is not exact but it is probably not far off. The probability that the difference does not exceed a certain amount can be easily calculated using statistical theory. It depends only on the true answer and the number of trials. We refer loosely to that difference which is just as likely as not to be exceeded as being the "error" in the calculation. For a solitaire game which has a probability of approximately 1/2 of success, and 1000 trials, this probable error would be about .01 or 2%. A number of trials larger than 1000 will give us a smaller error and vice versa. Of course, these trials could also be performed on the high speed machine instead of being played by humans. It is then quite easy to run a lot of them. Let us now consider one way of applying this technique to our problem.

Simulating the Area Defense

Looking at either Charts 26 or 27 we notice that there is another way in which the information may be presented. In the upper portion of these charts, we have drawn two lines. On the top line the tick marks are separated by exactly the height of the bars. We can measure the probability of any number of planes penetrating by comparing this upper line with the second line. This second line has a scale reading from 0 to 1 and can be used to measure the distances between tick marks.

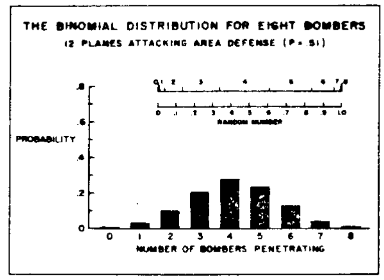

CHART 26

Chart 28 illustrates one kind of random number generator. In this case it is a wheel which can be used to produce numbers randomly and uniformly between zero and one. If the pointer is twirled it has the same probability of ending up between .1 and .11 as between .45 and .46, or in fact in any interval of size .01.

If now we produce a number from this random number wheel and mark this number off on the lower line of Chart 26, we can match up this number with a particular event specified by a certain number of planes penetrating. For example, if the random number wheel produced .2, we can take it to

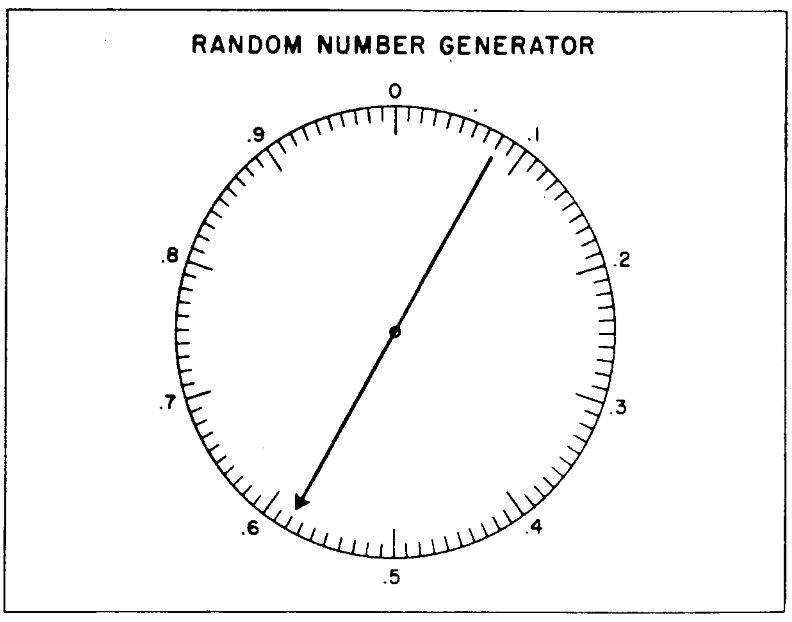

CHART 28

mean that 3 planes survived; .5 means 4 planes survived, and .9 means 6 planes survived. It is clear that as we go through this process many times, <u>we shall in fact simulate a situation in which the probability of getting any particular number of planes is given by the length of the corresponding bar</u>.

It is our purpose to use this method of producing random events to simulate the attack on the area and local defenses, keeping track of both bombers and escorts. At first sight one might think that we would need separate charts for every possible combination of escorts and bombers and different sets for the area and local defenses. If this were true, then if we wanted to consider force sizes of up to say 20 planes, we would need 400 charts for the area defenses and a similar number for the local defenses. This is a lot of charts. We can reduce the number greatly by displaying the information in a slightly better fashion.

In Chart 29, we have transferred the upper line of Chart 26 to a similar line. The places where the ticks used to be are now given by the points where the curves cross the line. In fact, the curves on Chart 29

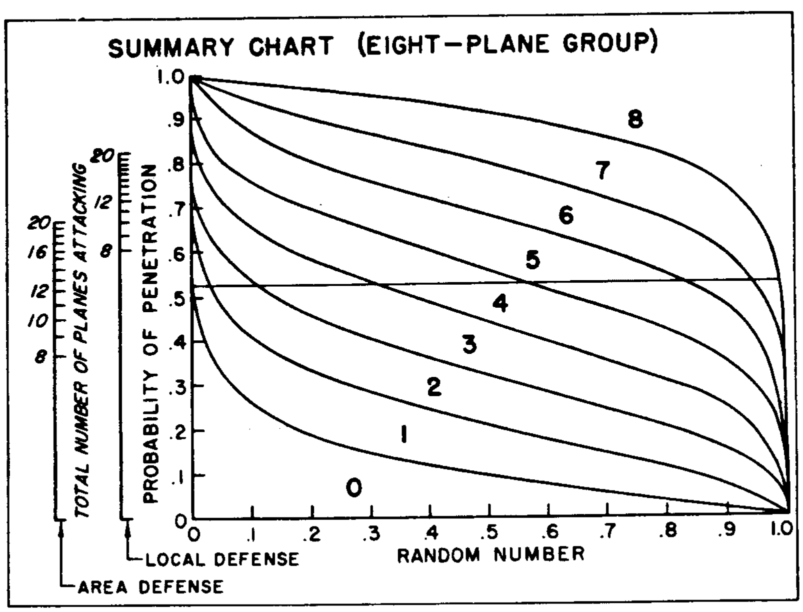

CHART 29

are drawn so that we can read off the probability of getting any particular number of planes penetrating out of an 8-plane group which is part of a larger force. The survival probabilities are, of course, determined by the total number of planes attacking and by whether these planes are attacking area or local defenses; and we have two axes labeled in this way. If we didn't do this we would have to refer to Charts 5 and 7 to get the survival

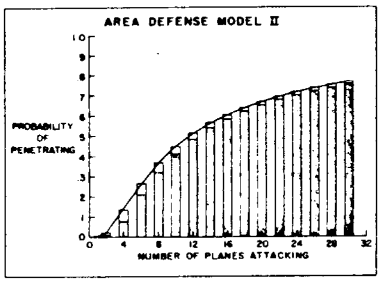

CHART 5

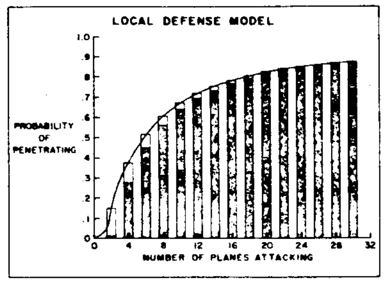

CHART 7

probability and then use the probability axis (on the right of the three vertical ones). Now, if we want to go as high as 20 planes, we need only 20 such graphs rather than the 800 that would be necessary if we presented the information as in Chart 26. The first 12 of these charts are given on pages 149 to 154 in Appendix 1.

<u>The Monte Carlo Calculation</u> (done in class by the students)

Let us now see in detail exactly how we would Monte Carlo this problem. Normally, we would turn the wheel every time we need a random number.[1] This time let us just list the first ten random numbers that could have been obtained

.10, .33, .77, .14, .35, .55, .81, .09, .39, .75, .38, .05, ...

Incidentally the RAND Corporation has put out a book which contains 500,000 independent two digit numbers. This book can be obtained for $10.00 and we don't recommend it for light reading, or any other kind. We should

---

[1]Actually we didn't have the wheel when we were writing so we looked them up in <u>A Million Random Digits</u>, The Free Press, p.77, line 10, or maybe it was p.92, line 38.

also mention that there are many standard jokes about these numbers. We are thinking of putting out a book of them too.

Let us continue, however, with our Monte Carlo calculation. Assume that we have an offensive budget of $100,000,000. Consider first a case where we buy 12 planes and 8 bombs (8 bombers and 4 escorts) and use a Type I operation.

MONTE CARLO WORK SHEET

	TRIALS / FORCE	1		2		3		4	
		AREA DEF.	LOCAL DEF.	AREA DEF.	LOCAL DEF.	AREA DEF.	LOCAL DEF.	AREA DEF.	LOCAL DEF.
BOMBERS	8	2	1	3	1	5	2	5	1
ESCORTS	4	2		2		1		2	
PLANES	12	4		5		6		7	
DEST.			.20		.20		.32		.20

CHART 30

We shall begin by simulating the attack against the area defense. The first random number is .10. It is used to determine what happens to the bombers. Looking at Chart 29 we find .10 on the abscissa and draw an (imaginary) vertical line. We draw another (imaginary) horizontal line from the ordinate corresponding to 12 planes attacking the area defense. The two lines cross in the area marked 2. We conclude that we have been kind of unlucky and that on this trial only 2 bombers have penetrated the area defense. We write a 2 in the proper space of Chart 30 and continue.

The next random number is .33. We shall use it to find out what happens to the 4 escorts. Looking at Chart 31 for the 4-plane group we draw

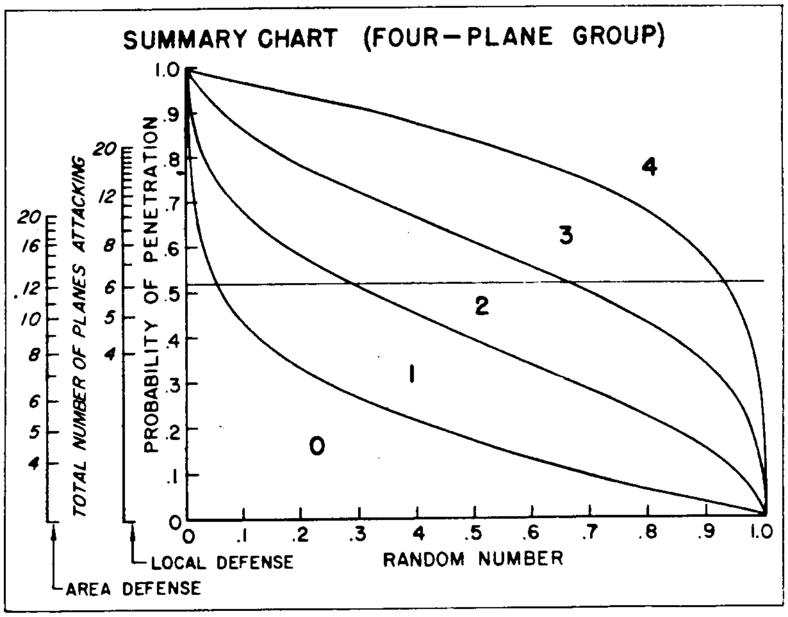

CHART 31

a vertical line from .33 and a horizontal line from 12 planes attacking area defense. These two lines also intersect in the 2 area. So we have 2 escorts penetrating with the 2 bombers.

We are interested, of course, in finding the number of bombers which penetrate the local defense. The cell which is attacking the local defense is composed of 4 planes but only 2 bombers. We are therefore interested in a 2-plane group. (See Appendix 1 for graph.) The random number is .77. Hence one of our bombers gets through and drops a bomb on target. This bomb destroys 40% of that airfield or 20% of the total target system.

We must run a large number of trials like this, starting always with 8 bombers and 4 escorts, before we can get an accurate picture of what is happening. Chart 30 shows the results of a relatively small number of such calculations. Averaging these results we get an estimate of .23 for the average destruction. A complete set of calculations gives .29 for the answer, so we seem to have been a little unlucky in our simulated attacks.[1]

Analysis of a Fixed Budget (Average Destruction)

It should be clear that we could do the same thing for a Type II operation and for other allocations. Chart 32 give the results of a complete set of such calculations for a $100,000,000 budget. In order to

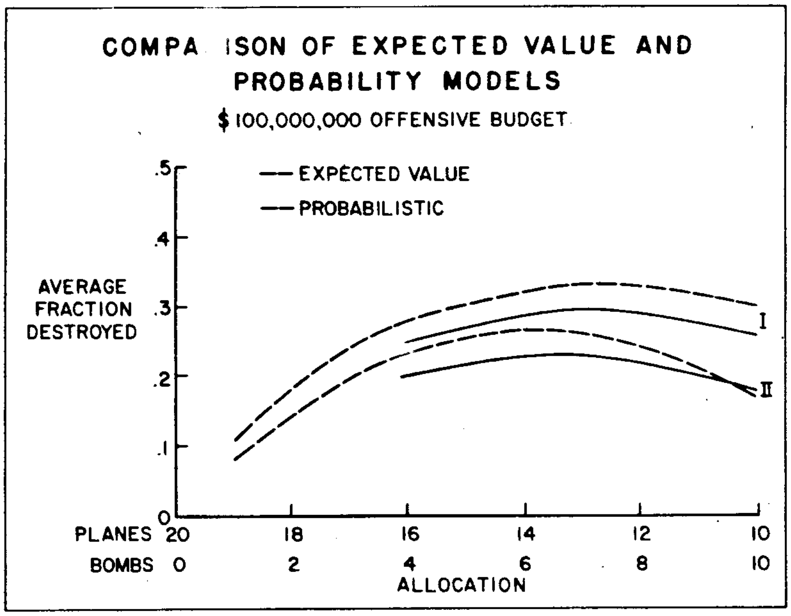

CHART 32

facilitate comparison we have also drawn on this chart the results of the expected-value model given by Chart 13. As can be seen, the expected-value model was optimistic and overestimated what we could accomplish, at least for most allocations. However, the error of the expected-value model is ordinarily somewhat less than the difference between operation types. The most significant thing shown by the chart is that an allocation which works well for the expected-value model will also work well for the

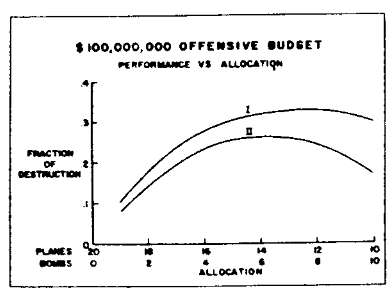

CHART 13

[1]The probable error for four trials is actually .05.

more accurate probabilistic model. This would occur even if the curves were not as flat as they are because the effectiveness of most allocations is changed by about the same amount by the introduction of probabilistic considerations. The flatness of the curves just reinforces this effect. Therefore, in this case, a study on allocation or operation type would not have to go to the probabilistic model. The expected-value model is good enough. It is not only good enough, it is about **one tenth the work**.

We are not, of course, implying that one can always trust expected-value models--only that this is a typical case of how preliminary study and thought may enable one to dispense with complicated techniques in favor of simpler ones. In practice, it is surprising how easily one can often reduce the amount of labor required by doing some preliminary analysis.

Results of Probabilistic Calculations

We have, of course, run through a probabilistic calculation for every possible budget and found the maxima. These are shown on Chart 33. This

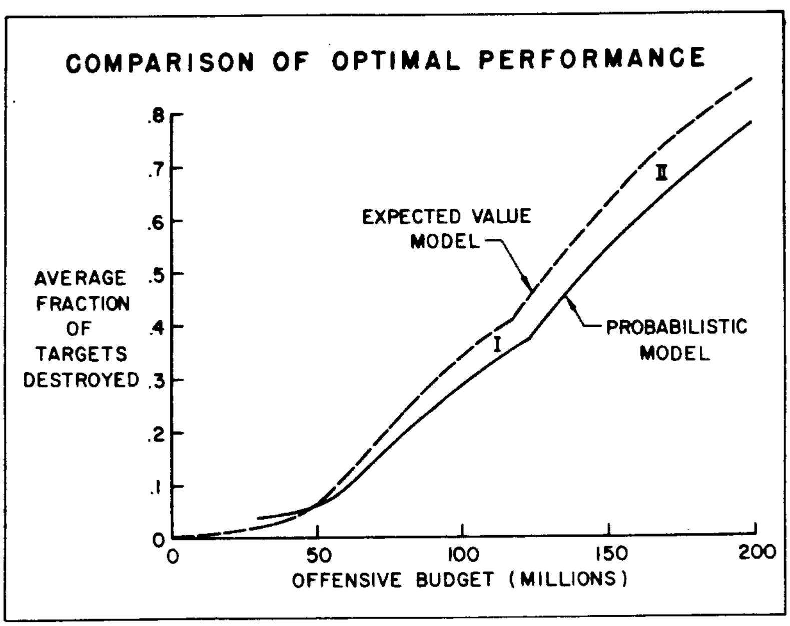

CHART 33

chart also has results of the expected-value model as given before on Chart 18. As we can see (and probably anticipated from the "S" shaped curve), when we are poor the expected-value model underestimates the efficiency of our forces; when we are rich it overestimates it. The error, while interesting, is for many purposes not significant. The inaccuracy in our model probably swamps it. There are situations though when, even with poor models, one might still be interested in an accurate calculation.

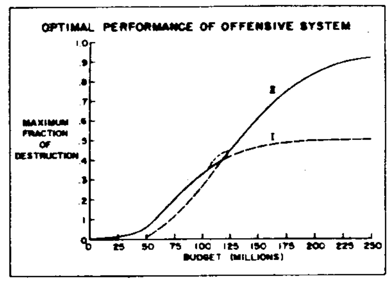

CHART 18

Deferred-Decision Operation

Once we begin thinking in probabilistic terms, we have opened up new alternatives in operating our air force. For example, the probabilistic calculation comes out a little badly not only because of the curvature of the curve, but also because we can no longer predict what the composition of our force will be after it has gone through the area defenses. If, for instance, we happen to be operating near the cross-over point between the Type I and Type II operations ($120,000,000 budget), then if we are lucky and don't lose many planes to the enemy's area defenses we should probably do a Type II operation; if we are unlucky, we should probably do a Type I operation. To give another case in which it is advantageous to make a decision rather late in the game, consider the problem of keeping the proper balance between the number of escorts and bombers. It is perfectly conceivable that, if one is doing a Type II operation, all the bombers in the cell sent to one of the airfields are shot down going through the area defenses and only the escorts are left. Under these circumstances, there is no point in having the escorts continue on to their target alone.

One would want either to have these escorts join the second cell, or to take some of the bombers from the second cell and put them in the first.

In general, we would expect to gain fairly substantially in the effectiveness of our force if we could defer the decision on both the tactic and the target assignment until after we have gone through the area defenses. Chart 34 shows the gains that can be achieved by doing this.

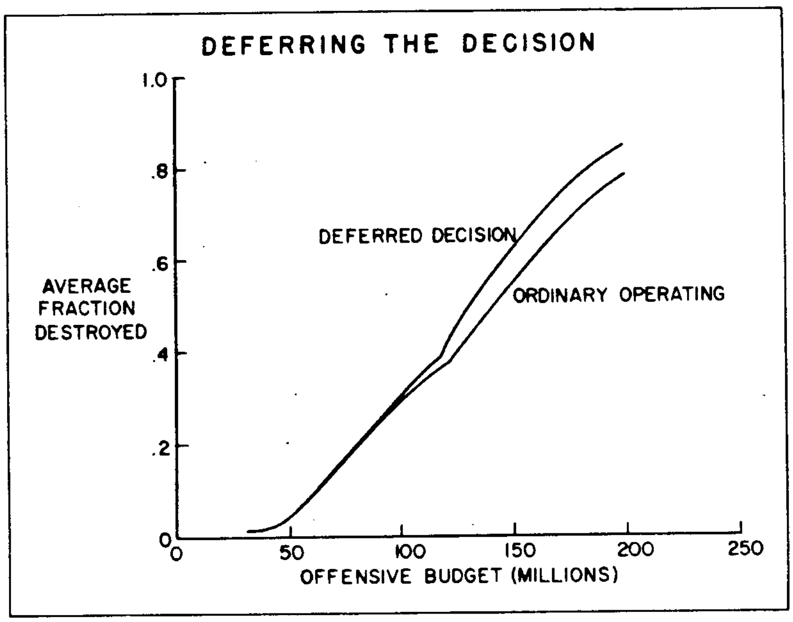

CHART 34

There is no gain when we are very poor because there is nothing we can do except send all of the planes which penetrate to the airfield we originally intended to attack. There are gains when we are rich enough either to be close to the cross-over point or to be operating Type II, because we can now plan our allocation on the basis of what actually happens in the area defenses rather than on what happens on the average.

Comparing Apples and Oranges

We sincerely hope that the reader has a slightly uncomfortable feeling at this moment. For one thing, we have ignored completely the relatively greater operational complexity of our so-called deferred-decision operation, so that we are not really maing a fair comparison between the two alternatives. In fact, it is probably true that the difficulty in doing this kind of regrouping over enemy territory is so great and the consequent degradation of our performance so large that it does not pay to try it.

We are definitely not making the point that one should disregard unorthodox strategies because they look operationally difficult--just the opposite. However, in making a comparison one should be aware that there may be costs or operational difficulties associated with the systems being compared that are not usually taken into account because they are subtle or hard to estimate. It is then usually best to use an obvious overestimate of both the unknown costs and the operational degradations of the preferred system. <u>If, after using this overestimate, one still prefers the same system, then one is in a position to make a recommendation</u>. We shall have more to say about this later when we talk about <u>a fortiori</u> arguments.

<u>Modelism</u>

Actually in talking about the deferred-decision operation, we have committed another sin besides comparing apples and oranges. We have pushed our model much too far. The whole concept of the deferred-decision operation implicitly assumed that the area defense was a thin barrier and that after one had passed this barrier there was time and space in which to regroup before reaching the local defense barrier. While the model suggests this way of thinking, actually, of course, both barriers occupy considerable amounts of space and there is no neutral area between them. Therefore, if one tries this reallocating strategy and waits until the planes penetrate fairly deeply before making a decision, one is likely to find that they are traveling very circuitous routes and are actually spending more time in the area defenses than if they had traveled direct routes. This, of course, substantially increases attrition. This is a typical trouble of many models. They may be quite satisfactory for answering one kind of question and yet misleading when used for other questions.

We should perhaps admit another sin at this point. Chart 34 shows in exaggerated degree the numerical value of the gain that is achieved by a deferred-decision operation. We exaggerated in this way because the actual gains were too small to show up effectively on the chart.

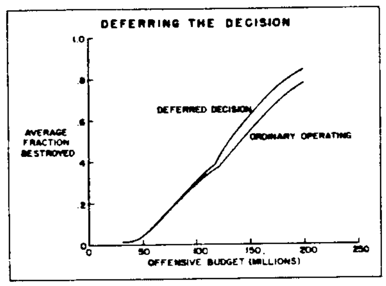

CHART 34

Probabilistic Objectives

As we have just seen, the probabilistic model gives a different and more accurate answer than the expected-value model. However, in our case, the results of the two models are close and the extra labor involved in Monte Carloing the problem is great. We would therefore probably not be justified, on accuracy grounds alone, in using the probabilistic model in preference to the expected-value model. There is, however, another and more important reason for putting probabilistic considerations explicitly into the analysis. The reason is that we want to state our objectives probabilistically.

Up to now, we have taken as our criterion of efficiency the average damage inflicted. This might be sensible if we were going to fight a series of wars, one after the other, and we wanted to maximize the total damage done in all of these wars. Modern war seems to be of a somewhat different character from the classical ones. Our major interest is probably in winning the next one and letting later ones take care of themselves. In particular, the first blow may be so overwhelmingly decisive that we may wish not to maximize the damage done on the average on this first blow, but wish instead to maximize our probability of achieving a satisfactory level of damage.

At first sight, these two objectives may not seem very different. As a matter of fact they often are. In particular what we do depends on what level of damage we consider satisfactory. If, for example, we try to operate in such a way as to maximize the probability of achieving say 30% damage, we will concentrate our attack on one field and lose completely our capability of achieving 60% damage. On the contrary, if we try to maximize our probability of achieving 60% damage, we must attack both fields. This may reduce appreciably our probability of achieving at least 30% damage.

Deterrence vs Win the War

Before discussing in detail exactly how incompatible the two objectives are, let us consider one reason why one might want to achieve 30% or 60% damage. It is generally recognized that the major objective of the armed services today is not to win the war if it breaks out but to prevent war from breaking out. One way to do this is to make it clear to the enemy that he cannot win the war. If we can convince the enemy that if he starts a war he will be destroyed, irrespective of whether or not he destroys us, then it presumably will not pay him to start a war. We say that he is deterred from starting a war. This deterrence mission has been publicly stated to be the primary mission of SAC.

Presumably one way to achieve this kind of deterrence would be for our forces to destroy a sufficiently large portion of his forces so that his later attacks will be too small to overwhelm us completely. This would correspond to having as an objective not maximizing our probability of winning the war handsomely but only maximizing the probability that he will not win the war handsomely. We shall assume that this last objective corresponds to achieving 30% damage.

The other objective, winning the war handsomely ourselves, would correspond to a 60% level of destruction. That is, we shall assume that if we destroy this percent of the enemy's forces, our active and passive defenses are sufficiently effective to cope adequately with whatever he can then send against us.

To repeat, we shall consider two different objectives: destroying 30% and 60% of his forces on our first strike. The 30% destruction corresponds to guaranteeing to the enemy that the most he can do from that point on is to get a technical victory. For example, two of him and one of us left. It would not be a satisfactory result even if the sexes were right.

The 60% level of destruction corresponds to a situation where we have hit him so hard that our own air defenses can more than take care of his attack. Most of us survive the war, and our society continues to go on in a more or less customary manner. It was almost a painless war.

Requirement for Deterrence and Win the War Objectives

Chart 35 shows the different kinds of events which are associated with different levels of destruction. For example, 1 bomb on 1 field and

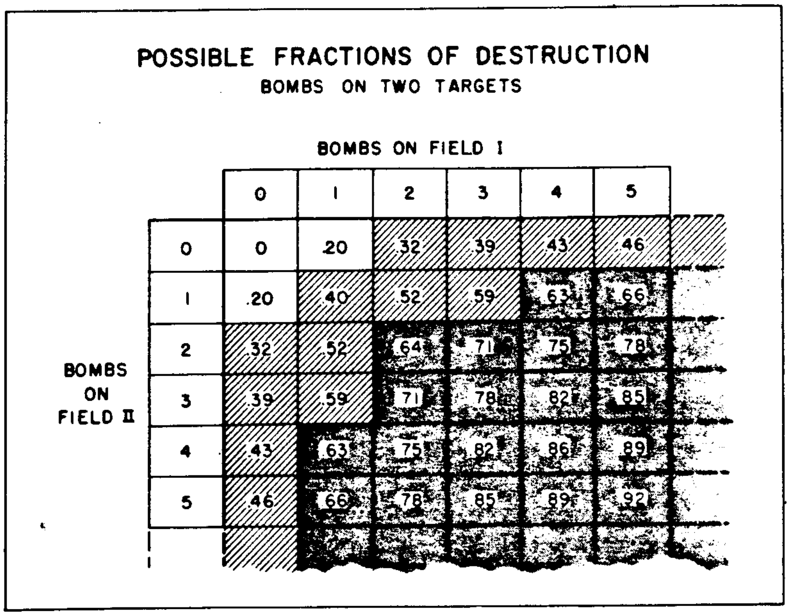

CHART 35

2 bombs on the other field will destroy 52% of the target system, 20% on the first field and 32% on the second. To take another example, 2 bombs on one field and 3 bombs on the other field will destroy 71%.

The unshaded area shows the events which correspond to fulfilling neither the 30% nor 60% objectives. The lightly-shaded area shows what kinds of events fulfill the 30% objective, and the heavily-shaded area shows what kinds of events fulfill the 60% objectives. You will notice that, if we are interested in a 30% objective, we give ourselves no extra credit for destroying more than 30%. In other words, every lightly-shaded box is taken as being as satisfactory as any other lightly-shaded box even though the boxes correspond to destructions that range from 32% to 59%. Another note of unrealism is reflected in the fact that we have included the 59% box in the 30% objective rather than stretching a point slightly and putting it in the 60% objective.

Sample Calculation of Probabilistic Objectives

Assume now that we have a force of 14 planes and 6 bombs, a typical $100,000,000 force. Chart 36 shows what this force could expect to do under a Type I tactic. It shows the probability, given this operation,

PROBABILITIES OF SINGLE EVENTS

14 PLANES, 6 BOMBS

B	0	1	2	3	4	5	6
P_B	.105	.247	.297	.219	.102	.027	.003

PROBABILITY OF ACHIEVING AT LEAST 30% DESTRUCTION: .648

CHART 36

of dropping any particular number of bombs on the field. The shaded area indicates the numbers of bombs which will fill the 30% objective. The total probability of doing this adds up to .648. We have a .352 probability of not achieving our 30% objective. We have, of course, no probability at all of doing 60% damage.

Chart 37 shows what happens when the same force uses a Type II operation. It gives the probability of obtaining any particular combination of events. For example, the probability of dropping one bomb on one airfield and another bomb on the other airfield is .123. Under a Type II operation, we do have a non-zero but small chance of achieving our 60% objective. These events correspond to the heavily shaded area. The

PROBABILITIES OF JOINT EVENTS
14 PLANES, 6 BOMBS

	P_B	.506	.351	.124	.019
$P_{\bar{B}}$	$\bar{B}$ \ B	0	1	2	3
.506	0	.256	.178	.063	.010
.351	1	.178	.123	.043	.007
.124	2	.063	.043	.015	.002
.019	3	.010	.007	.002	.000

PROBABILITY OF ACHIEVING AT LEAST 30% DESTRUCTION .388
PROBABILITY OF ACHIEVING AT LEAST 60% DESTRUCTION .019

CHART 37

probabilities add up to roughly one chance in fifty. A gambling man might be willing to take the chance.

Our probability of doing 30% destruction is appreciably better, adding up to .388 or roughly two chances in five. This value is, however, a sharp comedown from the .648 probability we could have achieved for our 30% objective if we had used a Type I operation. Of course, we might be willing to make this sacrifice in order to gain the one chance in fifty of a real success. We might figure, given that the war has broken out, that there is not much to be gained in mutual suicide. Unfortunately, in the real world one cannot make this decision after the war has broken out because the decision involves doing things which cannot be changed at the last moment. If one takes deterrence seriously, one might well plan on what is practically a <u>contingent</u> mutual suicide in the hope that the plan will <u>prevent</u> the contingency from occurring. (There is some additional discussion on this point in the chapter on Game Theory--in particular note the remarks on the rationality of irrationality.)

Analysis of a Fixed Budget (Probabilistic Objectives)

We have done this same calculation for every possible allocation of $100,000,000. The results are given in Chart 38. If one really wants to

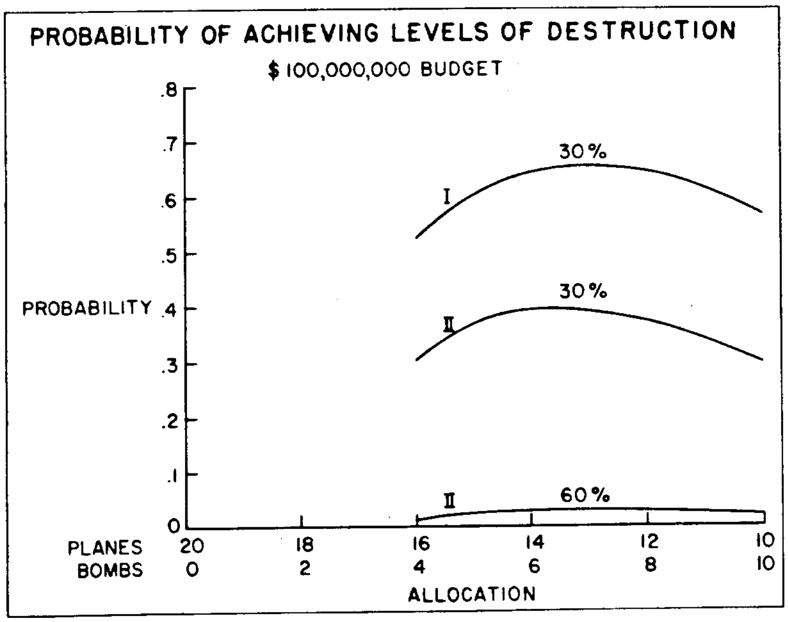

CHART 38

achieve 30% destruction he had just better use a Type I operation. If one insists on having a chance to see his grandchildren after the whole thing is over, he will have to operate Type II and correspondingly reduce his deterrence. There simply may be circumstances in which the enemy would not take a 65% chance of suicide but is desperate enough to take a 39% chance.[1]

[1]This is an important point about deterrence. Among other things one measures the value and strength of his deterrence by how desperate he can afford to make the enemy and not by the probability that the enemy will attack in cold blood when there is no tension. It is probably not necessary to have much of a deterrent force if one is never going to aggravate the enemy or frustrate his plans.

The second thing to notice is that the curves are still flat. The precise allocation between planes and bombs is not crucial. Forces whose allocation is optimized on the basis of the average destruction criteria, whether with the probabilistic or expected-value models, are still allocated accurately enough for our more sophisticated objective.

Performance Under Deterrence Objective

Chart 39 shows one set of results that can be obtained by going through

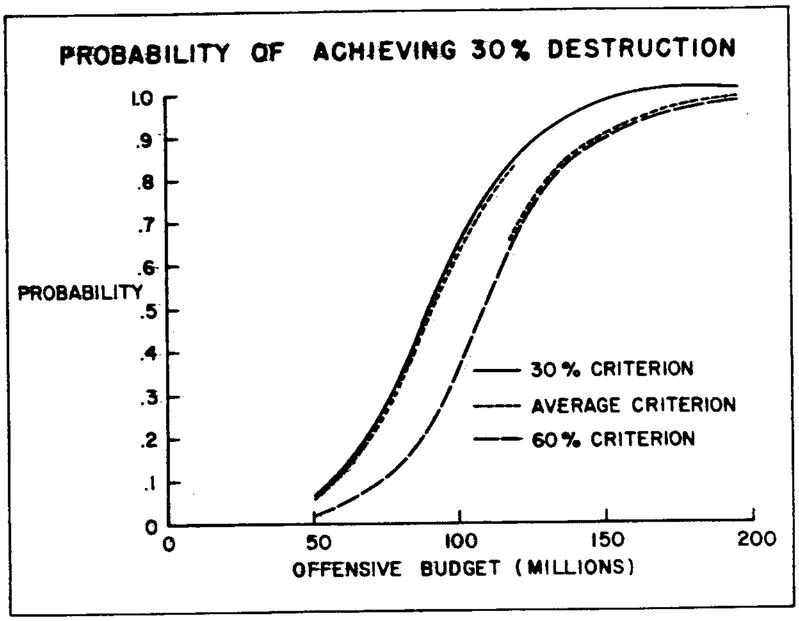

CHART 39

the probabilistic calculations for all possible budgets. The three curves give the probability of achieving 30% destruction, if the allocation and the operation are maximized under three different criteria: probability of 30% destruction, the average destruction, and the probability of 60% destruction. You will notice (as you perhaps expected) that there is essentially no difference between the average-value criterion and the other criteria when the operating type is the same. There is a large difference, however, between the 60% or 30% objectives and the expected-value objective, if it misleads one into using the wrong operation type.

We should not be fooled by the apparent closeness of the curves. The vertical or horizontal distances are better measures of the difference than the perpendicular distance. The vertical distance gives the difference in performance for a fixed budget; the horizontal distance gives the extra budget required for fixed performance; the perpendicular distance does not give anything.

It is also important to notice that if one is very rich and operates under the 60% criterion there is still a big probability of achieving 30% destruction. Even a 30% fanatic would then probably agree to the 60% philosophy. Likewise, if one is very poor, the possibility of operating successfully under the 60% criterion is so low that a 60% fanatic might give in to the 30% man.

There will probably be a real and intransigent controversy if the budget is in the center region between say 75 and 125 million dollars. Many people believe that such controversies can be objectively resolved by a more careful analysis of the issues. This may be so, but we are doubtful. In any case we are willing to bet that there will be very few multiplications in the analysis. However, we do believe that such research should be encouraged, even though we are not very hopeful of obtaining conclusive results.

Fundamentally, the choice between the two objectives is probably a matter of taste, and like most controversies about taste probably cannot be settled by objective quantitative logic. The relative values a country gives to these two different objectives reflect its ideology or value system. Ideologies and value systems are, currently at least, not very amenable to analysis, except possibly of the psychoanalytic variety. Therefore the attempt to decide the relative values of the two objectives seems to come under the heading of what we have called intuitive judgment rather than considered opinion.

Performance Under Other Objectives

In addition to calculating how well the 30% objective is satisfied if the optimization is done under the three criteria, one is interested in what can be done for the 60% and average destruction objectives under the different philosophies of operation. These are shown in Charts 40 and 41 respectively.

The two charts are self-explanatory. The main reason for showing them in addition to Chart 39 is to emphasize the necessity of considering multiple objectives and the compromises that must be faced when designing

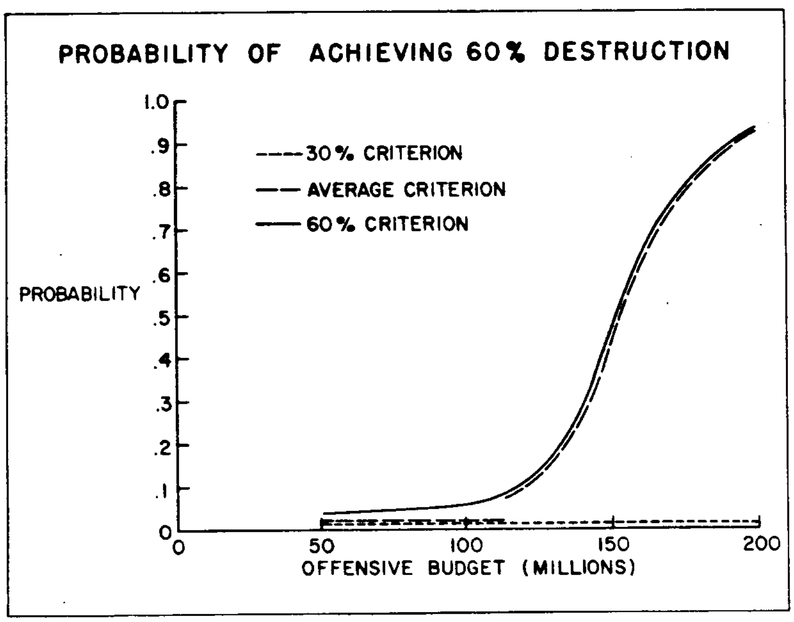

CHART 40

for them. In effect, the set of Charts 39-41 replaces the single Chart 18 or Chart 33 of a more naive study. They illustrate that one should consider possible contributions to all reasonable objectives in evaluating a system.

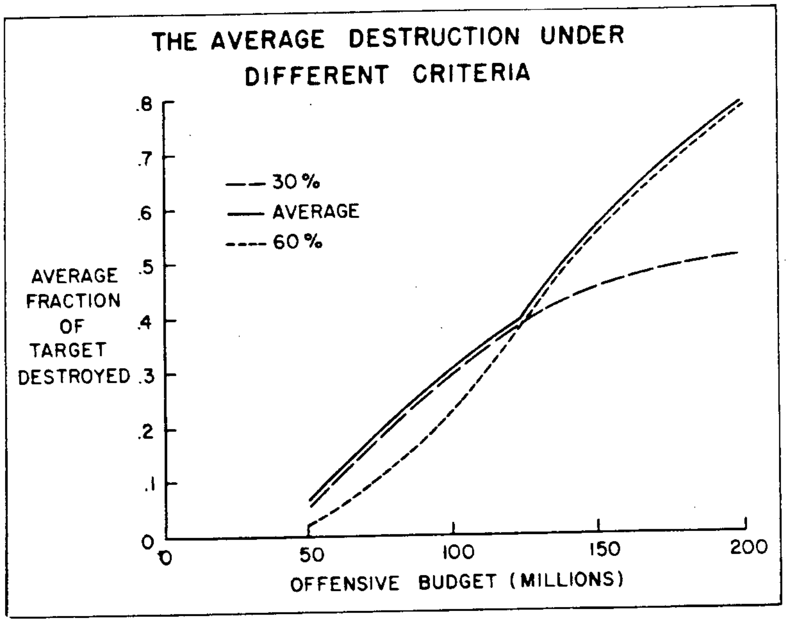

CHART 41

It is probably clear that the argument between proponents of different objectives is as much an internal controversy within an individual as between

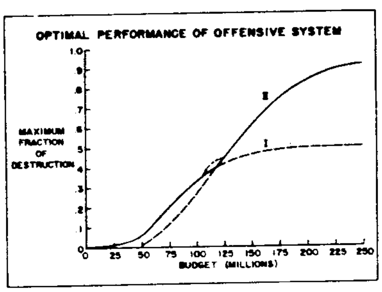

CHART 18

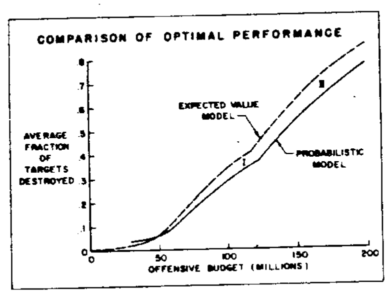

CHART 33

different people. We are all interested in deterring a war and, failing deterrence, winning it. We must, therefore, not only make our social compromises but our individual ones also. Some of the considerations that go into making such compromises are discussed in Chapter 3.

High and Low Confidence Measures

Let us go back to Chart 33. If one considered that chart seriously, one might make the statement that if we had less than $105,000,000 we would always fail to achieve our 30% objective while if we had more, we would always succeed.

RM-1229-1

The actual situation is not like that at all. Chart 39 shows that with $105,000,000 we have a probability of around .73 of achieving our objective. If we increase our budget we increase our probability, and vice versa. There is no sharp boundary line at any particular budget.

A 73% chance of achieving an objective is a pretty high probability of success but it might not be high enough for some purposes. Statisticians like to talk about the 90% or 95% confidence level as being roughly equivalent, for some purposes (but not all), to certainty. We can see from Chart 39 that these levels would require budgets of roughly $130,000,000 and $160,000,000 respectively. Probability levels as high as these we will refer to as High Confidence measures. They are not necessarily measures on which one can <u>always</u> place a complete and total reliance, but measures which are about as reliable as a reasonable person could expect to get.

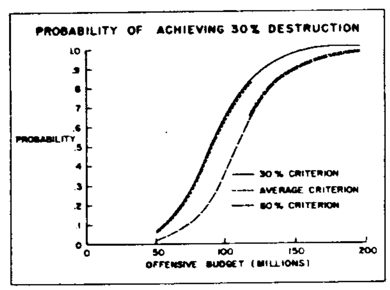

CHART 39

There is another class of measures of interest, which we call Low Confidence ones. Assume, for example, that we have a budget of only $60,000,000. Then if you turn back again to Chart 33, you will note that with $60,000,000, we can on-the-average get only about one-twentieth of the opposing target system. This would scarcely even annoy the enemy. The fortune looks useless. This view is not quite correct.

Turn now to the probabilistic model and look at Chart 39. You can see there that $60,000,000 gives better than one chance in ten of achieving the 30% objective; i.e., it has one chance in ten of preventing the enemy from winning the war satisfactorily. Therefore, our $60,000,000 has a certain amount of deterrence value. Even though the enemy has spent enough money to have, in a manner of speaking, a moral right to a satisfactory victory, there is a small but significant probability that he will either lose the war or not win it in a satisfactory fashion.

The issues are so big that even one chance in ten might make him hesitate. The contrast between the High and Low Confidence measures illustrates a fairly general point. The rich, having paid their money, just want the expected to happen. They are , in fact, willing to invest large parts of their fortunes just to decrease fluctuations. The poor, on the contrary, should lead a reckless life. It is only by great good luck that they can do anything at all. On the whole, their salvation lies in increasing the fluctuation, in making the situation chancy and uncertain.[1]

In most realistic analyses it would be misleading to talk about specific numerical probabilities as we have done here. One therefore often contents himself with talking only about Low and High Confidence measures, or sometimes Low, Medium, and High Confidence measures. For most purposes this three-point scale is satisfactory.

Expected-Value vs Probabilistic Calculations

Chart 42 lists the relative merits of expected-value and probabilistic models. We have shown (for our problem at least) that if one is interested in determining the best allocation of his money, or what operation type to use, or qualitatively how the performance varies, the expected-value model is likely to be good enough. If one wants to do more than this, one must go to the labor of a more accurate Monte Carlo or other probabilistic calculation. The

[1]There is a more careful treatment of this question in Chapter 10 on Game Theory (Part Two) when we consider games of ruin.

EXPECTED VALUE	VS PROBABILISTIC MODELS
1. ALLOCATION	1. QUANTITATIVE PERFORMANCE
2. OPERATION TYPE	2. LOW CONFIDENCE MEASURES
3. QUALITATIVE PERFORMANCE	3. PROBABILISTIC OBJECTIVES
	4. DETAIL

CHART 42

probabilistic models are not only more accurate, they also enable one to see how his system will react under good or bad luck. Last and most important, they enable one to talk about probabilistic objectives.

It is sometimes also said that one needs Monte Carlo in order to put detail into the problem. Whether this is true or not depends on the problem and what one is trying to do. However, it is almost always easier to put detailed considerations into Monte Carlo models. This is one of the reasons why, on map exercises, Monte Carlo methods are almost always used.

The Monte Carlo technique that we illustrated in the example was of the simplest and most naive sort. In Chapter 9 on Monte Carlo (Part Two), we discuss various kinds of improvements that can be made in this type of calculation.

CHAPTER 3: DESIGNING THE DEFENSE

We are now going to consider what happens when the enemy drops bombs on us--an unpleasant but necessary part of our study.

Bias for the Offense

We believe that it is reasonably correct to say that western military leaders (historically) have a bias for the offense. It is easy to see why this might be. For example, the Air Force probably wants its enlisted men and junior officers to be aggressive, audacious, and possibly even reckless. A pilot who is instructed to fly through an anti-aircraft barrage is in no position to compare the military worth of evasive action versus accuracy in dropping his bombs. If he is told to make this comparison in a cool calculating way, there is likely to be a great deal of highly emotional evasive action and very little straight line flying. In order to balance the reasonable regard that people have for the value of their own lives, military organizations train and select people to emphasize the offensive spirit. (Actually, of course, responsibility for other people's lives can be even more paralyzing. And, therefore, even leaders must be trained to have an offensive bias.)

Even though this offense-mindedness is desirable in operating people, it can have very regrettable results when it invades the planning and policy levels. It is often said that defense alone will not win a war. While this is patently true if taken in a ridiculous sense, wars have in fact been won and will be won in the future by emphasizing defense as much as, or in some sense more than, the offense.

We hear many warnings about Maginot-mindedness. It may be that most of these warnings are based on a misunderstanding of the lessons of history. We should first realize that the French built the Maginot Line as a reaction against the excesses that offensive minded generals committed in the first World War; secondly, that it was, probably, a good idea. If the French had built their Maginot Line without any holes, and used it properly to enable a relatively small number of soldiers to defend a large frontier, they could have concentrated troops where they needed them--to conduct offensive operations. It is one of the purposes of a good defense to enable one to pursue offensive operations when and where needed. A good defense not only prevents the enemy from destroying one's offensive capability; it causes him to divert and use up large resources in his offense. This presumably weakens his defense.

We don't want to belabor the point any more except to emphasize that the decision on how to allocate money between offense and defense should not be made on the basis of slogans but on the basis of how best to carry out the objectives we are interested in.

General Considerations

Let us now consider the problems of defending our strategic air force in a little more detail. We wish to allocate our money between area defenses, local defenses, and airfield shelters. Our objective will be to maximize the number of planes that survive the enemy's attack. That is, we shall evaluate our defense only by how much offensive force it protects. As we shall see in Chapter 4, there are some other considerations which we are ignoring. In line with our conclusions of Chapter 2, we shall assume that an expected-value model is good enough for the qualitative analysis. Therefore, we shall ignore probabilistic considerations.

Our initial problem will be a little more complicated than the corresponding one of Chapter 1 because we have to choose three variables rather than two. It is, however, of the same general character.

The Parametric Defense Model

We start by finding out what we can buy for our money. This is given by the three sets of curves in Charts 43, 44, and 45. Consider Chart 43 first.

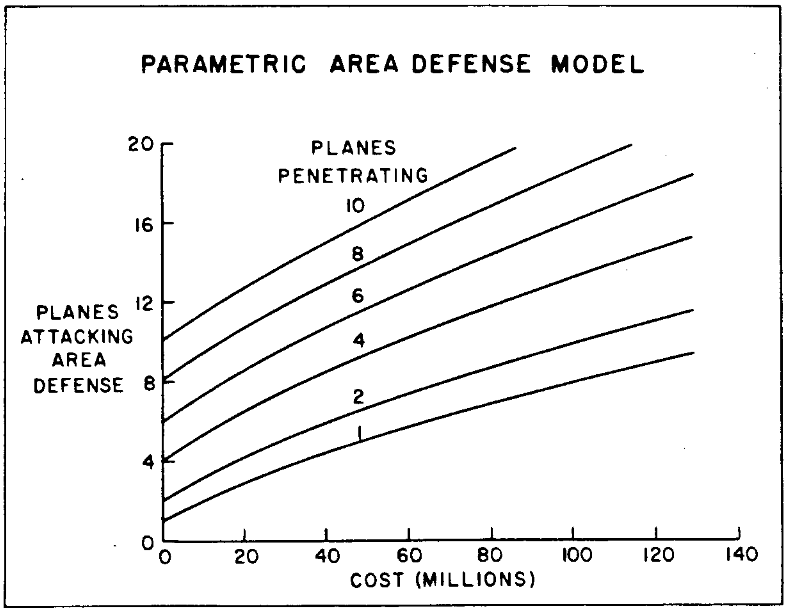

CHART 43

We have here a description of how effectively we can spend money to defend against different kinds of attacks. For example, if one draws a horizontal line from the ordinate that corresponds to 16 planes attacking, he will see that as we spend more money the enemy, on the average, gets fewer bombers through. If we spend $50,000,000, he will get an average of 9 1/2 planes through the area defenses; if we spend $100,000,000 he will get only 6.

We have a similar set of curves for the local defense model. The shapes and numerical values are slightly different but they are somewhat analogous to the area defenses. The third set, for the sheltered airfields, has a

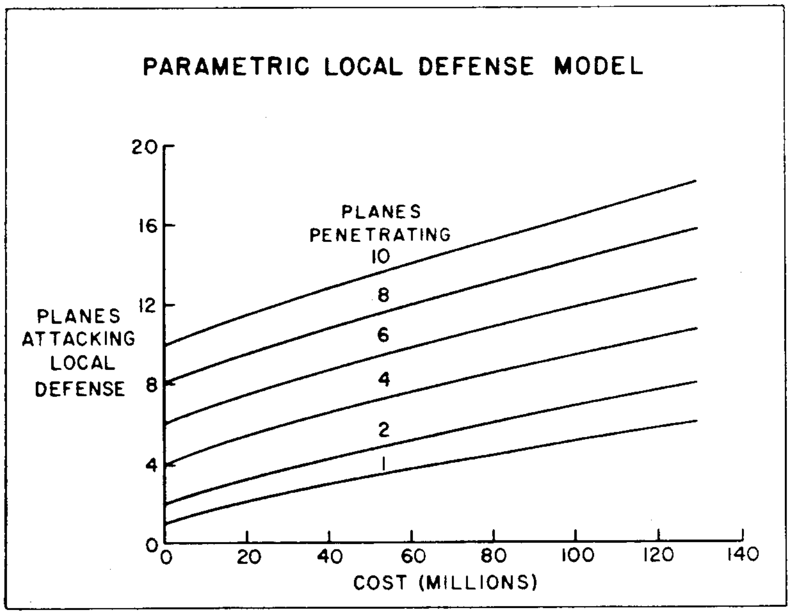

CHART 44

somewhat different analytical character, although not greatly different.

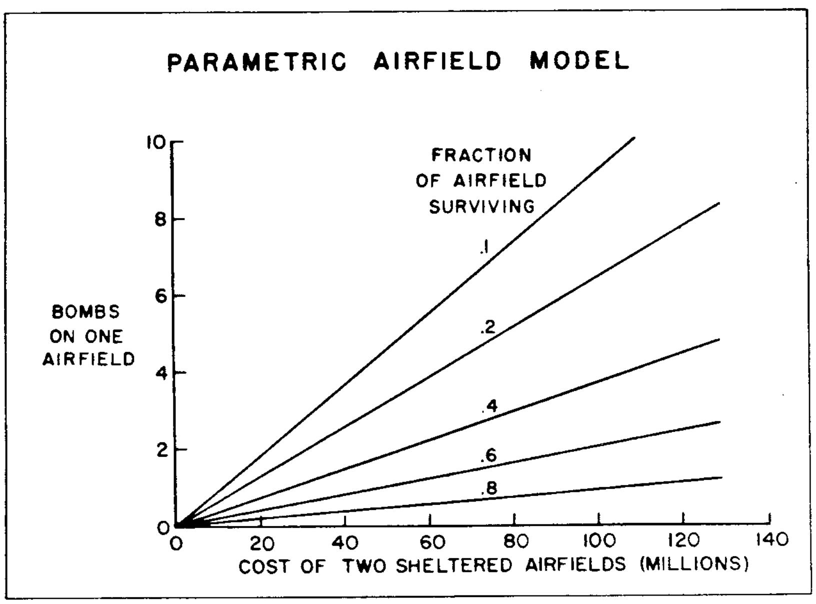

CHART 45

Now, where do we get charts like this? We just drew them. This is practically the classical way. More prosaically, they can be obtained by a combination of more or less detailed map exercises and (after a lot of prodding, pleading and cajoling)[1] preliminary engineering designs. They are almost never obtained from detailed studies. It would be too expensive to have a lot of detail and still try to cover a large range of parameters. In this kind of study one puts in this type of detail later, not when he is looking at a whole range of possibilities, but after preliminary choices of the interesting parameters have been made. However it is essential to have the rough curves to help make the preliminary choices.

This is not to say that one should not look at some details early in the study--only that it is not necessary to spend too much time on relatively well understood details. (There is more discussion of this point in the Ten Common Pitfalls chapter.)

The above charts then apply to a very specific type of situation and would have to be changed, perhaps drastically, if the situation changed.

Optimal Allocation of a Fixed Budget (another class exercise)

Assume that we have $150,000,000 to spend on defense and wish to allocate this sum in the best possible fashion. Assume also that our intelligence has told us that we can expect an attack of 8 bombers and 8 escorts (16 planes carrying 8 bombs). Let us start by fixing the local defense at some arbitrary figure, say $75,000,000. We can now divide the rest of the money between the area defenses and shelters. A series of such divisions is shown

[1]Engineers are quite accustomed to using this kind of very rough data in the preliminary stages of a problem. However, they often dislike giving the same kind of rough data to Systems Analysts for fear that it will be abused. It probably will be, but it is better to have a little abuse than to have a worthwhile study held up. There is more discussion of how to obtain and treat technical estimates in the chapters on Ten Common Pitfalls and Miscellaneous Comments (Part Three).

in Chart 46. We also show on this chart what happens when the enemy's force attacks our various allocations. We let him try both a Type I and Type II operation to see which is better for him.

WORKSHEET FOR DEFENSE ALLOCATION

DEFENSE BUDGET = 150 M | PLANES ATTACKING = 16 | BOMBS = 8

COST OF AREA DEFENSE	COST OF LOCAL DEFENSE	COST OF THE TWO SHELTERED AIRFIELDS	BOMBERS PENETRATING AREA	BOMBERS PENETRATING LOCAL		PROB. OF PENETRATING		BOMBS ON TARGET		FRACTION SURVIVING	
				I	II	I	II	I	II	I	II
0	75	75	16	11	7.5	.69	.47	5.5	3.8	.58	.54
12.5	75	62.5	14	9	6	.56	.38	4.5	3.0	.58	.55
25	75	50	12.5	8	5	.50	.31	4.0	2.5	.57	.54
37.5	75	37.5	11	6.5	4	.41	.25	3.3	2.0	.56	.52
50	75	25	10	5.5	3	.39	.19	2.8	1.5	.53	.49
62.5	75	12.5	9	4.5	2	.28	.13	2.3	1.0	.52	.39

CHART 46

The results of these calculations are graphed in Chart 47. The Type II

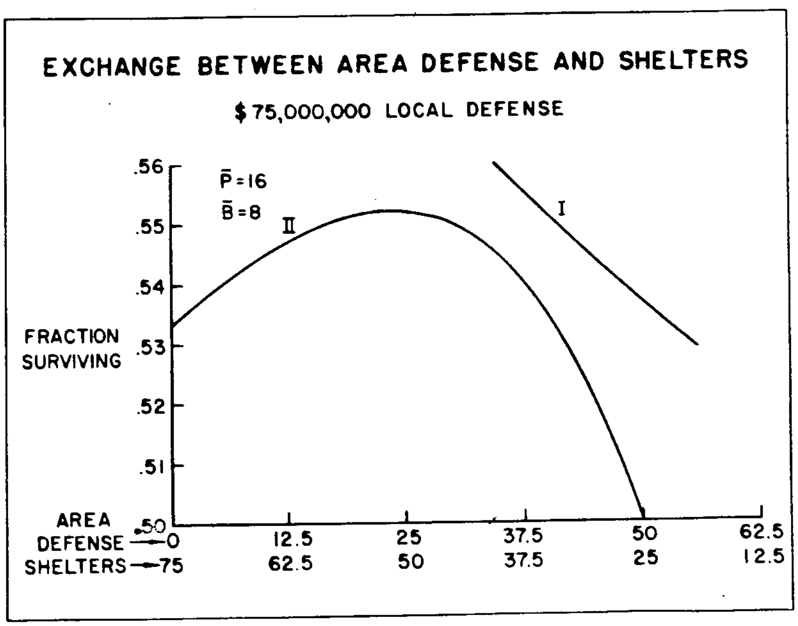

CHART 47

operation is always preferable from the viewpoint of the offense (survival is less than on a Type I operation). Therefore, the enemy should choose a Type II operation. We, of course, wish to make the survival as high as possible, and so we should pick the allocation that puts us at the maximum point.[1] This maximum occurs when roughly $23,000,000 is allocated to the area defenses and $52,000,000 for the shelters.

We continue with the calculation by fixing one of these two figures, say the area defenses ($23,000,000), and trade the rest of the $150,000,000 between local defenses and shelters. The results of the new calculation are shown in Chart 48. Once again we see that the enemy prefers a Type II

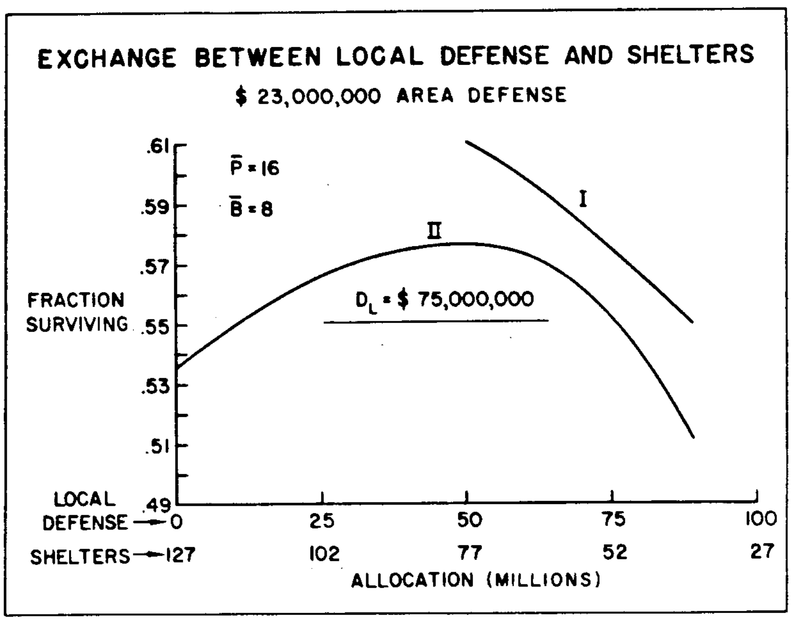

CHART 48

operation. We try to maximize our survival and choose $46,000,000 for local defenses and $81,000,000 for the shelters. The horizontal line shows where

---

[1] In the Game Theory chapter (Part Two), this is referred to more technically as picking the maximum of the minimum curve, or as the max-min strategy.

the maximum of Chart 47 is when transferred to Chart 48. The reason it appears to be in a different place is that we have changed the scale.

Fixing now the amount to be spent on local defenses ($46,000,000), we can vary the allocation of the remaining $104,000,000 between area defenses and shelters. We then get Chart 49. We haven't bothered this time even to

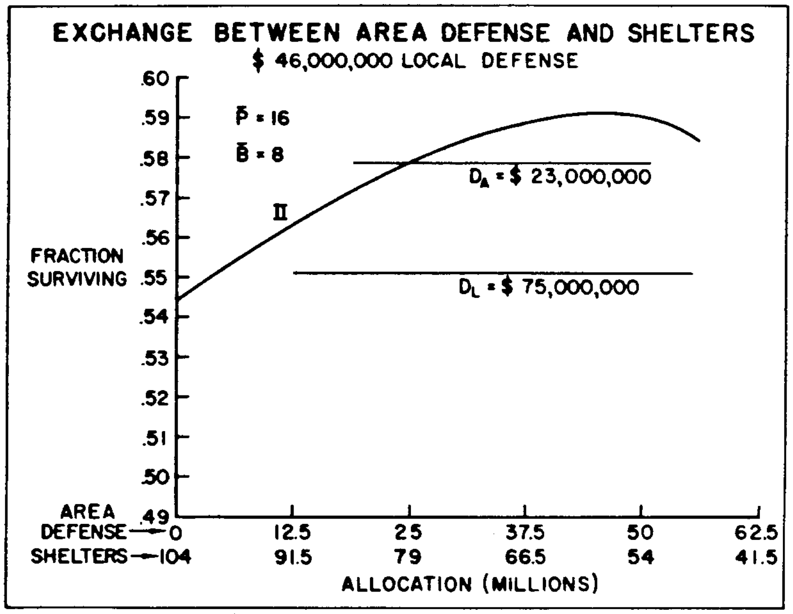

CHART 49

look at a Type I attack because we are convinced by now that the enemy should never use it. The optimal allocation comes out $47,000,000 for area defenses and $57,000,000 for shelters.

Fixing the $47,000,000 for area defenses, we trade the rest of the money between shelters and local defenses and get Chart 50. We still get an improvement but the improvement is quite small. The process can be continued indefinitely until the improvement is no longer worth the amount of work it takes. However, if we continuedthe calculation we would end up with roughly $50,000,000 for the area defenses, $40,000,000 for the local defenses and $60,000,000 for the shelters and get a survival of .53. These numbers are slightly different

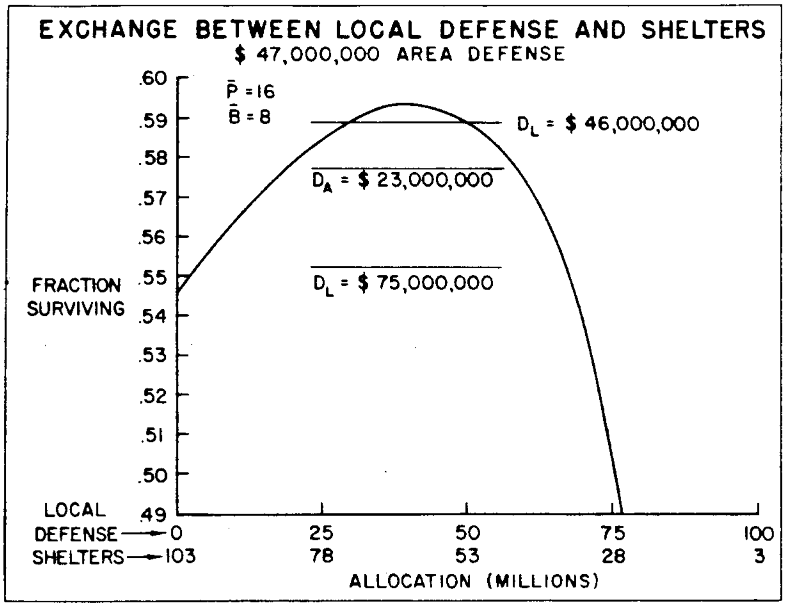

CHART 50

from the numbers shown in Chart 50 ($47,000,000 for the area defenses, $40,000,000 for the local and $63,000,000 for the shelters), but the performance of the two systems is just about the same. As always in a well-designed system the performance is insensitive to the exact allocation.

Maximizing Procedures in General

Let us now try to see in a general way what we were doing when we found the optimal allocation. Chart 51 is a kind of contour map which gives us our performance against a fixed enemy attack for any allocation of a $150,000,000 budget. We can look up on the ordinate the amount of money spent on area defenses; on the abscissa, the amount of money spent on local defenses; what is left over is spent on shelters. Our performance is given by interpolating between the two nearest contours at the intersection of any ordinate and abscissa.

In effect, Chart 51 is a contour map of a hill. The peak of the hill is the optimal performance point and any point off the peak is less than

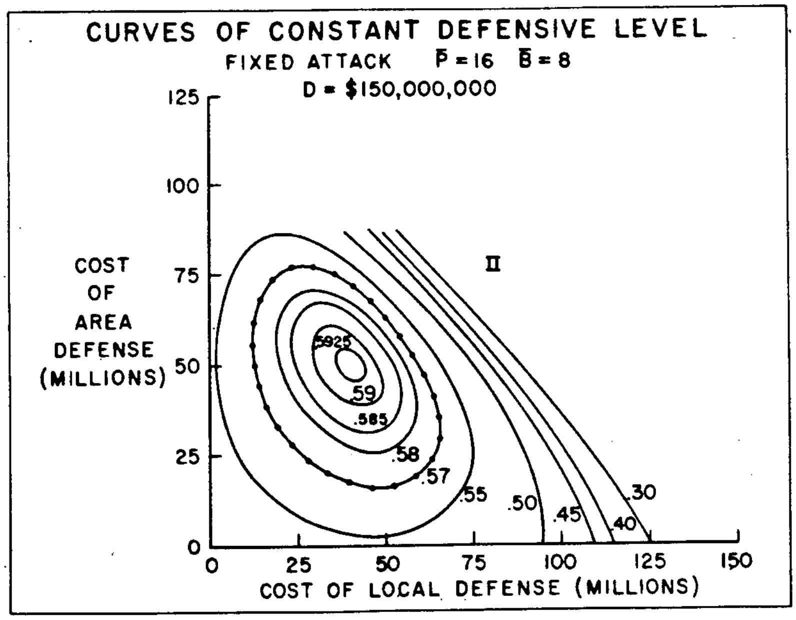

CHART 51

optimal. If we now review the previous optimization process on this contour map, we get the following. First we fixed the amount of money spent on local defenses at $75,000,000. We then walked due north and south along this line until we came to a high point. This occurs when the north-south line is parallel to the nearest contour and corresponds to a $23,000,000 area defense. We then walked due east and west on this area defense line until we again came to a high point, which occurred when the local defense was $46,000,000. We then walked due north and south on this fixed local defense line until we came to the high point at an area defense of $47,000,000, and so on.

The process can be repeated indefinitely until we arrive at the top of the hill. The first question one might want to ask is, "Is this an effective method of finding a maximum of a function of a large number of variables?" If one is working with graphs in the way we are and there are not too many hills, it is probably not a bad way at all. However, if one is working with calculating machines, particularly with high-speed computers, it is probably quite poor. There is another method which is much better, the so-called "method of steepest ascents." This method goes as follows: one picks a point at random and finds the direction in which the hill is going up most steeply. He then points himself in that direction and starts walking. There are two variations of interest. In variation one, he continues walking in the same direction until he comes to a high point. He then stops, finds again the direction of steepest ascent and starts walking in that direction. In modification two, instead of walking in the same direction until he comes to a high point, he goes a fixed preassigned distance and then looks around for a new steepest ascent direction.

Both of these variations are quite good and can often be used to find the maximum fairly easily. There are, unfortunately, cases where the contour map is very twisted, full of gullies and ravines of an indescribably horrible shape. Finding a maximum can then be very tedious. However, we are generally only interested in fairly broad, flat and smooth maxima. These are the ones which are insensitive to specific assumptions and are usually easy to find. We don't really care about the very narrow tortuous maxima even if they are quite high.

The main reason for discussing this problem of finding a maximum of a function of the three variables (area defense, local defense, and shelters), aside from its intrinsic interest, is that we want to show how the problem increases in complexity as the number of variables increases. In fact, it doesn't take very many variables before we find that we have exceeded the capacity of even a high-speed computer. In this case, one must use approximate maxima. We will give an illustration of how this is done in Part IV.

Contingency Planning

When we did this sort of thing before, we implied that this is a pretty poor analysis. It is. One reason is that we have not yet made any attempt to take account of changing circumstances and uncertainties in the other assumptions.

Uncertainty in Objectives

In Chapter 2 of this study, when we compared the 30% and 60% wars, we were trying to take account of uncertainty in our objectives. Doing this complicated our lives to some extent. In fact, it raised some unanswerable problems--unanswerable in the sense that logic alone cannot resolve them. In such a case, we said that it was essentially a matter of taste or of national policy or intuitive judgment. However, we may have made too much ado about the uncertainty in objectives. It turns out typically that this uncertainty is one of the lesser of the uncertainties with which we are afflicted, at least when we are talking about a large scale central war. (In a limited war, of course, it can be crucial--for example, the decision to use or not to use atomic weapons will probably depend on the political objectives of the war.)

Uncertainty in Costs

There is, also, the obvious counterpart of objectives, costs. There are almost always large uncertainties here. We have set ourselves the problem of designing a system that will achieve a certain set of objectives at minimum cost or, what is essentially the same thing, which will maximize our capabilities given a fixed level of cost. (This is sometimes loosely phrased as a desire to maximize capabilities for minimum cost. It is easy to show that this statement, if taken literally, is inconsistent or nonsense.)[1]

---

[1]The chapter on Elementary Economics (Part Two) will have more to say about this.

At first sight, costs look like a pretty concrete thing. You just grab an accountant and put him to work. However, as always, a close examination makes the precise imprecise. As we will explain later, when we talk about time-phasing, there is a real ambiguity in deciding the dollar cost of a system. Roughly speaking, this occurs because a military system is not bought at one instant of time by going into a department store and ordering it. It has to be built up over the past years and it is expected to have a continuous existence in the future. Under such circumstances one must always ask himself what it costs to use facilities which are already owned, and what will be the salvage value of any expenditures made this year in future years. Also, if one is procuring or developing a new system, he may have had no experience on which to base cost estimates. It is surprising, in practice, how inaccurate even careful estimates of the costs of new systems have proved to be. Careless estimates tend to be out of this world.

There is another ambiguity in costs which the Analyst generally ignores but with which the policy maker is sometimes concerned. Some dollars are harder to come by than others. Research and Development funds, for example, are ordinarily tighter than procurement funds. In the U. S., expensive gadgets are often easier to buy than high grade, but relatively low cost, people. Public works funds, for obvious reasons, are often easy to get and, of course, traditional expenditures are almost always easier to justify than new ones.

Finally, there are many costs which are not usually measured in dollars, such as crew lives, dislocations, political effects, etc.

Context Uncertainties

It turns out that the cost uncertainties can be quite large. One of the traditional ways to slant a study is to do the costing judiciously.

However, (aside from actual fraud) the uncertainties in cost tend to be like the uncertainties in objectives--annoying but not devastating. The really crucial uncertainties are those associated with trying to predict what the environment will be during the lifetime of the system being studied. While it would be wrong to design a system that works best under circumstances which are very improbable, it is more certain than either death or taxes that some of the circumstances that will actually occur will be improbable. There are so many ways in which improbable events can occur that some are sure to happen. Nothing is more improbable than the absence of improbable events.

We must, therefore, do more than just take into account the existence of a complicated set of objectives and different methods of doing cost accounting. We must also design our system in such a way that it will operate well under a large variety of circumstances. We call designing a system in this way Contingency Planning.[1] Some of the things which have to be taken into account in Contingency Planning are given in Chart 52.

1. Magnitude and Type of Attack

The first is the magnitude of the attack. Will the enemy have 1000 H-bombs, or only 50, ten years from now? Will their yield be 100 megatons or 100 kilotons? How many bombers, radars, fighters will he have? What kind of surprises?

You want a system which can exploit the situation, if he is weak, and yet not depend on his being weak.

---

[1]What we call Contingency Analysis some readers may be accustomed to calling Sensitivity Analysis. We find it useful to differentiate between large gross changes in context and smaller, more continuous changes. We have reserved the name Sensitivity Analysis for the study of these smaller changes and use Contingency Analysis for the large changes. We will not do any explicit Sensitivity Analysis in the example.

SIX IMPORTANT CONTINGENCIES

1. MAGNITUDE OF ATTACK
2. DEGRADATION OF OUR PERFORMANCE
3. VARIATION OF ENEMY PERFORMANCE
4. CHANCE FACTORS
5. STRATEGIC AND POLITICAL CHANGES
6. TECHNOLOGICAL PROGRESS

CHART 52

2. Degradation of our Performance

The second thing we might want to consider is the degradation of our estimated performance. Notice we said degradation, not variation.

The reader is undoubtedly familiar with examples of this effect. For instance one might ask a porter what his average tip is. He answers $2.00. You give him $2.00 and he says, "First time anybody came up to my average."

It should be understood that most of the estimates of average performance are of this sort. They are goals, often idealized and optimistic goals, rather than sober predictions. We are talking here about more than the (very important) question of the degradation of performance of individual soldiers and organizations. The Systems Analyst is equally concerned with the degradation of performance of equipment undergoing Research and Development as compared to the predictions of the designers. It is amazing how uniformly optimistic most of the contractors and official sources seem to be--at least for the short term estimates.[1] For the long term the same sources tend systematically to underestimate long term improvements. We shall discuss this latter bias under point 6 on Technological Progress.

[1]It should be noted though that there seems to be a sharp difference in the biases when the estimates are obtained from sources competing for funds as opposed

(footnote continued on next page)

There is another aspect of the degradation of our performance which is of extreme importance. It is related to or identical with the "battleship thinking" of the prewar Navy. The battleship was a fine object in its day but, when its day had passed, some of our naval officers were reluctant to give it up. They placed a reliance on it which proved to be costly.

Let us give a homely example. Assume you are in charge of defending a town which is surrounded by a swamp and that this swamp is traversed by a single road. If the road is fortified, you can credit yourself with a fairly effective system. The enemy is faced with the choice of either attacking the fort or accepting the risks and costs associated with crossing the swamp.

Assume now that some civic minded idiots build another road through the swamp. You would now be foolish indeed to think that you have a reliable defense simply because the main road is fortified. Yet history has shown that you will almost certainly have an enormous reluctance to give up the confidence that you used to have. After all, the fort still has a reassuring appearance of great strength. It stands there as massive and omnipotent-looking as ever. It is hard to believe that it is now almost useless (at least by itself).

You might think that this analogy is too trivial to be taken seriously. Unfortunately, it is literally true that most defense systems simply fail to plug up holes anywhere nearly as fast as they develop, even when the holes have been clearly telegraphed.

3. Variation in Enemy's Performance

Similarly, we must take into account the variation in the enemy's performance. In this case, it is really variation; there are systematic tendencies

---

to sources that have a comfortable (possibly monopolistic) position. These latter are more "responsible." That is, they hate to stick their neck out. As a result they are often overly cautious in estimating current capabilities.

both to underestimate and to overestimate. For an example of the first, we need only recall the many remarks to the effect that the Japanese could not see straight and therefore couldn't aim. We all know they fought very creditably. Similarly, we know of many people who think of the Russians as a collection of peasants or mujhiks, appalled by the mysteries of the internal combustion engine. Some of these people have been very unpleasantly shaken by the recent Russian fly-by's and weapons tests, but others have been scarcely stirred from their complacency.

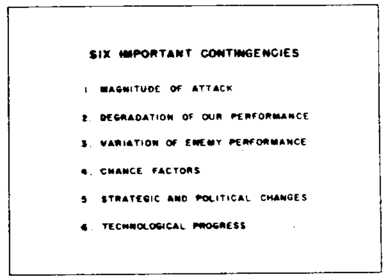

CHART 52

The other side of the coin is adopted by people who assume the enemy is omnipotent. One often finds an analysis which, in effect, lets the enemy put most of his gross national product into airplanes. Another analysis has the enemy spending the same money on atomic energy, and a third spends it on an expansion of his defense networks. By ignoring the many claims the enemy has on his resources, these analyses have assumed the enemy can eat his cake and have it.

In addition to giving him resource constraints, we must also consider the possibility of exploiting his stupidities. We have referred to a situation in which we had the responsibility of defending a town surrounded by a swamp. The town was defenseless because we had fortified only the main road through the swamp and not the auxiliary road. Now it would be premature to give up the fort just because there was a weak spot in the defensive system. First one can (and should) build a second fort to defend the back road. Furthermore even if the second fort is not built it is still possible that the enemy will attack along the fortified road. He may say to himself, "You put a fort there. You must have a reason. Let's attack it."

As a matter of fact we may have been a little unfair in our last comment. The attacker is, for several reasons, more likely to think through his alternatives exhaustively than the defender is. However, even though deception and finesse in the attack are common, they are not universal. One should, therefore, be willing to invest the small amount it takes to maintain the fort, just on the chance that the enemy will be stupid enough to hit you where you are strongest. This, however, does not excuse one from bestirring himself to defend the back road as soon as possible.

4. Chance Factors

There are really two kinds of soft spots. One kind just depends on the enemy's choice as to whether he wants to exploit it or not. In this case, it should be counted as a real weakness even though one may not think that the enemy is clever enough to find it. Sometimes though, the enemy's ability to exploit a soft spot depends on chance rather than on his free will. One might then debit himself less for the soft spot although not too much less, unless the probabilities are very low. For example, one might have a good capability for air defense in fair weather and a poor capability in bad weather. We can give ourselves more credit for having a good weather capability than we could if the enemy had the choice of the weather under which he attacks. This does not imply that we will not be extremely worried about the bad weather capability and not try to redesign our system so as to improve it. It's just that we will not trade as much of this good weather capability to increase our bad weather capability as we would if the enemy could choose the weather.

The principle still holds even when the enemy can make a choice, if the choice is costly to him. For example, our daytime air defense might be superior to our nighttime air defense. At first sight, this does not seem to do much good; the enemy can choose to attack at night. However, it may be very costly or inconvenient for him to defer his attack. If we can estimate this cost, we can give ourselves some credit for forcing him to come in at night.

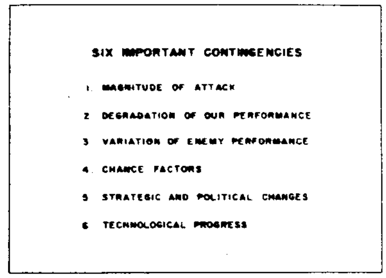

CHART 52

5. Strategic and Political Changes

One of the characteristics of our world is that hegemonies, alliances, and neutral areas continually change. We must keep our system sufficiently flexible to meet these changes. This means, for example, that we cannot rely on any particular set of foreign bases, though presumably we can place a fairly high reliance on having at least a portion of them available to us. It also means that we have to worry, for example, about the Russians taking over or influencing some countries which we have counted on as being allies or neutrals. We should not therefore make ourselves too dependent upon any particular group unless we have reason to believe that they are just about as reliable as one of the forty-eight states. (We presume that we can depend on Texas.)

Other kinds of strategic changes may be important. For example, it is very possible that future wars, like Korea, will include all kinds of target or other limitations, particularly as to the number or yield of atomic weapons that can be used.

Three important problems arise in considering such limitations. First, one must have a capability for enforcing the limitation on the enemy. In particular we must preserve our massive retaliatory capabilities from being destroyed by a sudden blow. Secondly, we must have a satisfactory capability for fighting all of the kinds of wars that may be necessary with whatever limitations or lack of limitations are involved. A third (and often overlooked point) is that we should operate as much as possible in such a way that the enemy is not tempted to violate limitations. It is sometimes possible to design a system to use effectively political, military, and psychological barriers.

In addition to designing systems that operate effectively within limitations, in principle we can consider refusing to accept limitations which would seriously hamper us. However, the questions of what limitations we are willing to accept is presumably not within the Systems Analyst's or even the Department of Defense's authority to decide. It is decided by either the President or Congress and the enemy. In some circumstances our allies will also have more than a slight say in the matter. It is clear that preserving a capability for fighting a whole spectrum of limited wars can be difficult--sometimes almost intolerably so. (It is important, however, that the Department of Defense planners realize that it is their job, as much as possible, not to hamper the civilians in their conduct of political and foreign affairs; military solutions which constrain the civilian arm of government in unaccustomed and serious ways should be looked upon as desperate expedients.)

6. Technological Progress

Finally, we must consider the effects of technological progress. Some of the most crucial uncertainties that bedevil a Systems Analysis come in here. It is at least partly because these are so hard to predict that the scholarly type has become useful on the policy level.

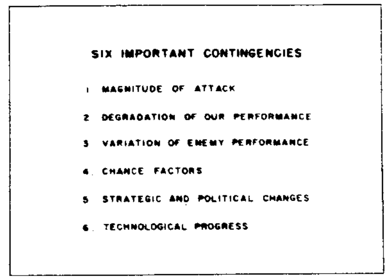

CHART 52

We think that, if experienced men are available whose experience has been obtained in competitive situations such that the incompetent are almost automatically eliminated, it is best to rely on their experience rather than on analytical techniques and analytical people. It is generally believed and probably correct that professors tend to be rather poor business men, even when their field is business administration.

However, it is characteristic of the current world that the effective rate of technological progress has been enormously increased. In a real sense, on many of the problems faced by the Department of Defense nobody has relevant practical experience. Furthermore, the need for good advance planning is also more important. Even in small wars, the pace of events may be so fast and the lead-time for development and procurement of forces or even for making operating changes so long, that it is practically impossible to correct mistakes in planning after the war has broken out. This is, of course, even more likely to be true in large wars. In this situation, one must somehow wed whatever experience is available to analysis. This means making scholars out of some of the military and military out of some of the scholars.

The history of the B-36[1] is a slightly atypical but not extreme example

[1] We are indebted to Albert Wohlstetter for suggesting this example. He discusses it more completely in his lecture in the course.

of how difficult it is to prepare for an uncertain future. It was designed during World War II when people were thinking first of Germany and then of Japan as the enemy. It was designed to carry high explosives. It was designed when its chief enemy was thought to be the propeller-driven interceptor.

None of the analyses which went into it and determined how we should trade range, weight, altitude and speed considered the possibility that:

- it might carry atomic bombs[1]
- the enemy might be Russia
- it would have to fight its way through jet fighters and guided ground-to-air missiles
- we would have overseas bases
- refueling techniques would be available

Any one of these changes might have been sufficient either to eliminate its value completely or to increase it enormously. Somehow, it is up to the man who is designing such vehicles to produce equipment which will be able to fight effectively in almost unpredictable situations.

In addition to proper design, there is one very important thing which can be done to alleviate this particular problem--to defer decision. One shouldn't decide today whether he wants to have a long-range slow airplane or a short-range high speed one in 1965. He should carry both projects through the paper design stages. If he still doesn't know in a couple of years which he needs, then he might carry both projects through the mock-up and possibly even the tooling stages. While the cost of doing this may be high, it is measured in millions and not in billions. It is therefore small compared to

---

[1]To give an illustration of how great this single change is we need only point out that one B-36 can carry a single bomb that has more explosive energy in it than the total of <u>all</u> the bombs carried by airplanes of both sides in World War II.

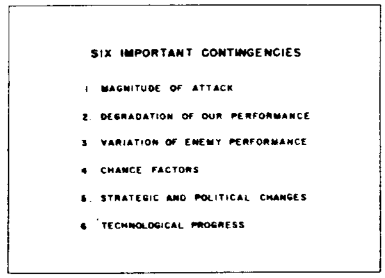

CHART 52

the total cost of the strategic air force.

One should always remember that the total investment in an organization like SAC runs between fifty and one hundred billion dollars. Therefore, in principle it would be sensible to spend even a few billion dollars to increase the effectiveness of this force by only 10%. Under these circumstances, one can clearly afford some very expensive development and preliminary tooling programs if they enable you to defer making crucial decisions until you can make them wisely.

We often hear statements that the major reason for doing a System Analysis is that development programs are so expensive and it is crucial that none of them be wasted; therefore, all development programs should be tied into a system, designed as a whole. Nothing could be more wrong. Development people have a saying, "It may or may not be a mistake to develop something which is not procured, but it is always a mistake to procure something which is not developed."

The most important thing the Department of Defense can do is see to it that we maintain a great deal of flexibility in our capability and have available a great variety of on-the-shelf items to meet a variety of contingencies. This ordinarily means that many of our development projects will never reach the stage of large-scale procurement. This may create very difficult relations with both Congress and the public. The problem has to be faced directly and preferably adroitly, but it is a mistake to cut back on potentially valuable development programs just to prevent possible criticism in the event they do not turn out to be needed. Nobody should object to buying insurance even if he doesn't have a fire.

There are many examples of how development programs conceived as a system can go wrong. An interesting one is discussed in a _Time_ magazine article on ICBM (January 30, 1956). _Time_ asserts (paraphrased):

> the ICBM program suffered from a lack of support because the guidance problems were so severe that the rest of the program was not pushed. The _unexpected_ development of the H-bomb suddenly made even very inaccurate ICBM's useful. We were unfortunately at that time in no position to benefit immediately from this development. We would have been in an even worse position but, luckily, an _entirely different program_--the rocket booster for the Navaho cruise type missile--had been pushed so far that we could use it as a basis for the ICBM engine.

This example is almost as spectacular as the history of the B-36.

It is also important to realize that even quite competent engineers and technical people can be very unreliable when it comes to estimating either the short or long range performance of systems undergoing Research and Development. This seems to be particularly true in the fields of electronics, aeronautics, and nuclear engineering. In the past 20 years there has been an exponential increase in the state of the art of all of these fields. Both as a cause and an effect of this exponential increase, these same fields play a central role in current military technology. For this reason military technology is much more unstable that most civilian technology and as a result the two very common mistakes mentioned occur. (Even first rate engineers will overestimate their ability to apply the new ideas in the immediate future and at the same time underestimate the rate at which long term development is proceeding.) To give a recent unclassified example from the electronics industry, let us consider the high-speed computer.

In 1945, 1946, and 1947, a large number of very competent engineers were going around promising to have wonderful electronic computers available in a year or two. In spite of these promises, by 1949, only computers of very modest attainments (as compared to the promises) had been built and the slippage of the more ambitious machines had become a joke. For example some of the engineers were saying, "We are really very reliable. No matter when you ask us we will always reply that the computer will be ready in 18 months." However, by 1950, a number of these better machines had achieved semi-reliable operation and, by 1951, they were almost all working well. Today, only 5 years later, there are machines on the shelf whose performance exceeds even the wilder and most futuristic extrapolations made in 1946. Furthermore, computers which will soon be available are an order of magnitude better than the current ones.

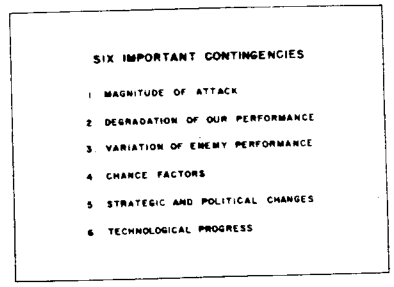

CHART 52

The above story illustrates the two most typical biases in technical estimates for rapidly changing fields. Because these biases exist the Systems Analyst should recieve technical estimates with some skepticism if not cynicism (and thereby often further embitter already embittered relationships).

Partly because of this technological uncertainty it is almost impossible to play in complete detail development programs for complicated systems like the B-36 or ICBM. This is even more true when we are talking about our long term military posture. The best that can be done is to push the state of the art in a whole series of component fields, give support and encouragement to competent people who have ideas they want to try out, be on the alert to extract by-product and bonus values and, most important of all, examine the programs as a whole to see if they are complete. They Systems Analyst should especially concentrate on the last two things. After all, almost everyone else is tied up with either specific projects, administration, budget crises, or congressional investigations. In some cases, he is just about the only Indian who can spend full time looking at the broader aspects of a program. What is also important, he often has a full-time and technically competent group of associates to help him look.

Simulating Different Contingencies

In our model the first four contingencies are essentially taken into account by just varying the magnitude of the attack. For example, the degradation of our performance looks analytically like a larger enemy attack. Similarly, the variation of the enemy's performance looks like a variable enemy attack and, lastly, enemy good luck can be simulated by giving him extra forces, bad luck by making his forces weaker. This method of simulating contingencies is not only useful in our model, but, to more than a slight extent, in the real world. Probably the single most effective gimmick in doing Contingency Planning is to consider a range of enemy attacks. One must of course add to this a detailed discussion of specific contingencies. However, in our little study, we shall, except for a few minor addenda, limit ourselves to the first.

Non-Contingency Planning

Chart 53 shows how not to do contingency planning. It does, we admit, look at three different contingencies, a small, medium, and large attack. It

NON—CONTINGENCY PLANNING
$150,000,000 DEFENSIVE BUDGET

CONTINGENCY ATTACK	P	B	D_A	D_L	D_S	SURV.
SMALL	12	8	72	25	53	.695
MEDIUM	16	14	50	40	60	.399
LARGE	24	16	0	0	150	.264

CHART 53

shows how best to allocate the defense budget if any of these attacks occurs and what the performance will be with any allocation and its <u>corresponding</u> attack. The first time we saw a chart like this, it was called contingency planning, but this is definitely a misnomer.

Contingency Analysis

The reason is shown on Chart 54, which gives the contingency analysis of the planning of Chart 53. It shows, for example, that, if one plans for a heavy attack and a small one actually materializes, he could have done 35% better if he had planned correctly. Similarly, if one plans for the small attack and the large one materializes, he could have done 67% better.

CONTINGENCY ANALYSIS
$150,000,000 DEFENSIVE BUDGET

CONTINGENCY MET / CONTINGENCY PLANNED	SMALL	MEDIUM	LARGE
SMALL	.695 (0)	.391 (2)	.158 (67)
MEDIUM	.681 (2)	.399 (0)	.195 (35)
LARGE	.513 (35)	.311 (28)	.284 (0)

CHART 54

Contingency Design

While looking at a chart like this is a step in the right direction, it is not enough. It is very common in discussions of Systems Analysis to have people come up with such a chart and then ask, "What do you do? How do you choose between these systems?" This attitude is more than a little wrong.

It implies that one has a rather small range of choices and that the big job is somehow to decide among these choices. This is essentially impossible to do satisfactorily in this extreme case because people have different estimates of what the circumstances are likely to be.[1] The main job of a good Systems Analyst is to design a system that will be satisfactory in all reasonable contingencies. He should spend most of his time trying to make the decision problem less agonizing rather than on the decision problem itself. In fact, it is fair to say that a good way to measure success in designing a system is on how unagonizing one has made this choice problem.

Let us see what this would mean in our analysis. Going back to Chart 51, we have marked extra heavily a line corresponding to a survival of .57. If one operates anywhere within the area limited by this .57 line, he will do at least 95% as well as the maximum (.593) that could have been possible. Losing only 5% is not very much. We shall adopt the attitude that, if we operate anywhere within this area, it will be reasonably satisfactory. Because we have deliberately designed the system to be insensitive, the area is fairly large. If it is not, we should go to a great deal of trouble to make it larger.

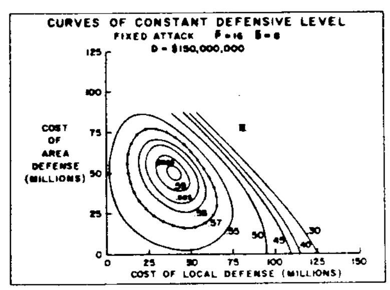

CHART 51

Let us now consider how we can exploit this idea when planning for the three attacks presented in the Non-Contingency Chart. We have drawn contour maps for all three of these attacks and they are given in Charts 55, 56, and 57. As before, we have marked heavily those lines which correspond to

[1]Chapter 8 (Part Two) on Probability and Statistics discusses some of the logical difficulties that arise when one tries to get a consensus for the probability of an event.

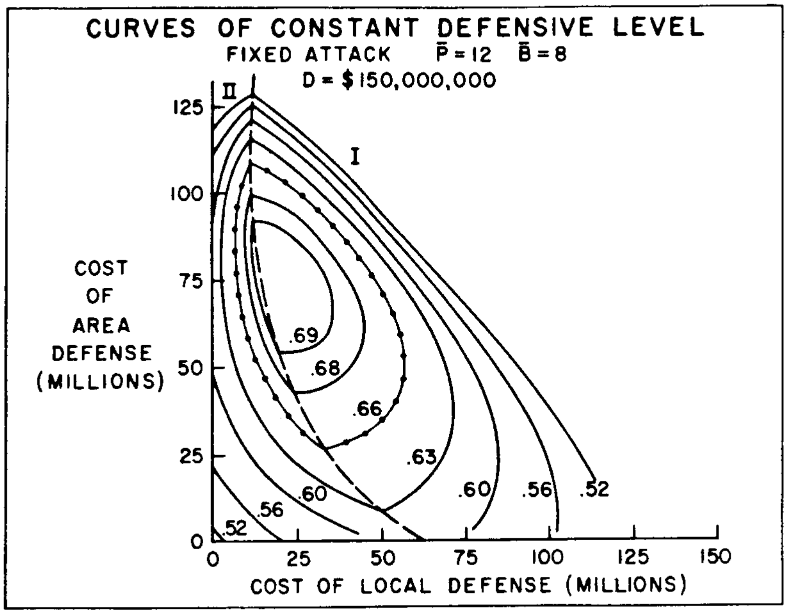

CHART 55

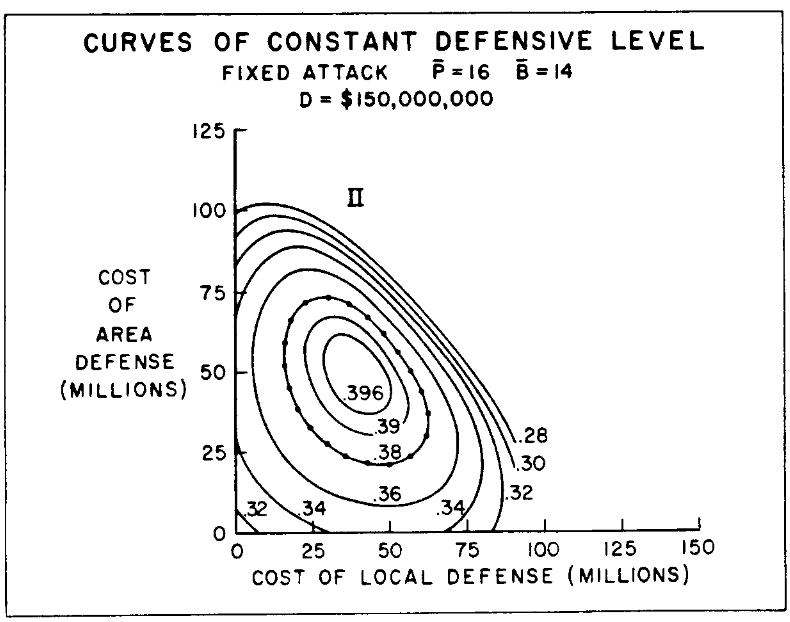

CHART 56

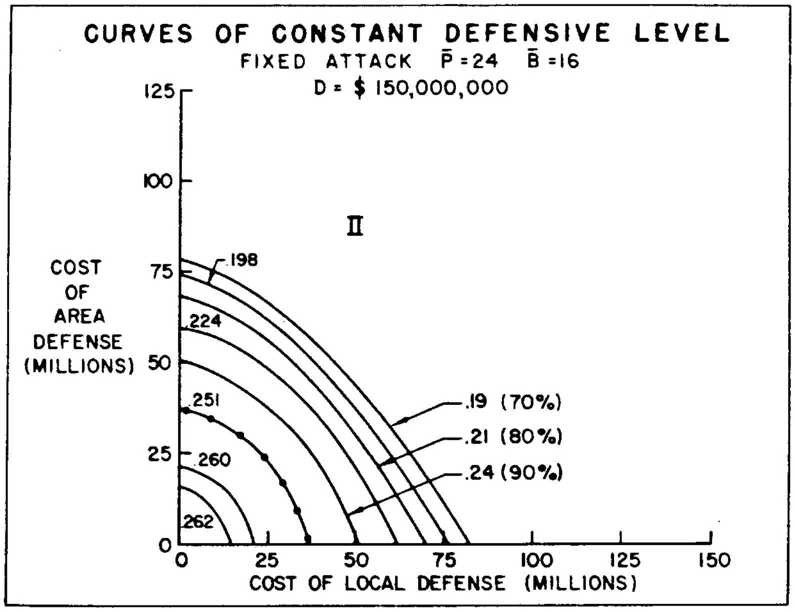

CHART 57

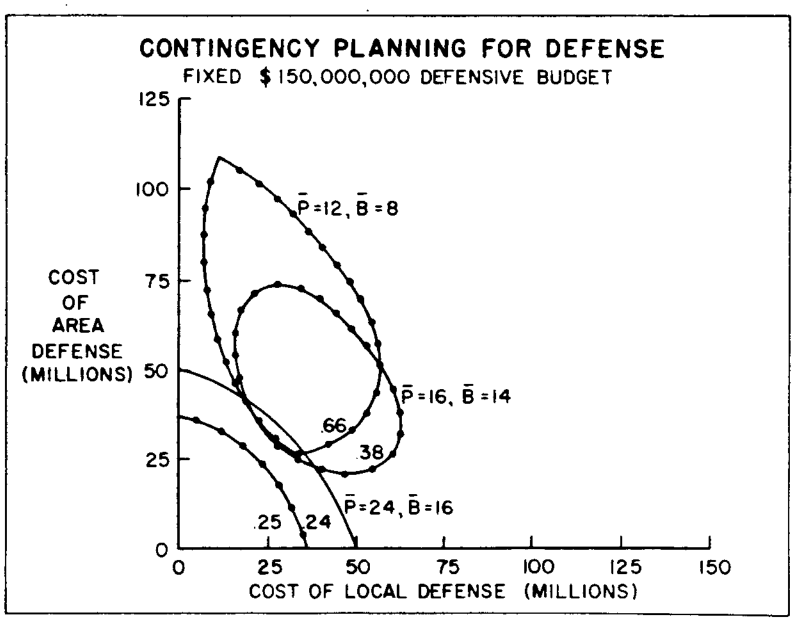

CHART 58

areas within which we do 95% as well as the best possible. We then transferred these 95% areas to a common chart (Chart 58). You will notice that the areas corresponding to the light and medium attacks overlap quite satisfactorily so there is no trouble in picking a compromise to meet these two attacks. Unfortunately, the area corresponding to the heavy attack does not overlap. Our first reaction should be to put some more work into designing the system to see if we can make it less sensitive, even at the cost of a slight loss in effectiveness. If this cannot be done, then we simply have to give up on our criteria of doing at least 95% as well as the best possible. After all 95% is not a sacred number.

It is a little hard to decide what should give. We might argue that if the attack is small or medium we are going to do pretty well anyway. We are really most worried about the desperate circumstances when the attack is heavy, so maybe we should relax on the lighter circumstances. On the other side, we might argue that actually the heavy attack is very improbable. It was picked, as a matter of fact, as being substantially heavier than we thought of as being reasonable. It would be a pity to compromise our capabilities for the probable event to take account of what is, after all, an extremely improbable event.

It is most unlikely that we will find an objective way to decide this argument. At this point therefore the analyst must leave the field of considered opinion for the intuitive judgment. Many of the people in military establishments seem (properly, we think) to prefer worrying about the more desperate circumstances. In order to give a little emphasis to the other side (but with no intention of giving Pollyannas any encouragement) we chose to yield the other way and drew a curve within which one does at least 90% as well as the best possible against the heavy attack.

There is now a reasonable overlapping area. We can pick a system anywhere in this area as our best compromise design. We did this and Chart 59

CONTINGENCY PLANNING A
($150,000,000 DEFENSIVE BUDGET)

D_A = \$ 30,000,000
D_L = \$ 30,000,000
D_S = \$ 90,000,000

CONTINGENCY ATTACK	PLANES P	BOMBS B	SURV. S	LOSS	PERCENT
SMALL	12	8	.662	.033	5
MEDIUM	16	14	.382	.017	4
LARGE	24	16	.236	.028	12

CHART 59

shows how our compromise system performs under the three attacks. It shows how much is lost in survival as compared with the best possible system specifically designed for that attack.

We think that everybody would agree that the system described in Chart 59 is better than any one of the systems described in Charts 53 and 54.

NON-CONTINGENCY PLANNING
$150,000,000 DEFENSIVE BUDGET

CONTINGENCY ATTACK	P	B	D_A	D_L	D_S	SURV.
SMALL	12	8	72	25	53	695
MEDIUM	16	14	50	40	60	399
LARGE	24	16	0	0	150	.264

CHART 53

CONTINGENCY ANALYSIS
$150,000,000 DEFENSIVE BUDGET

CONTINGENCY MET / CONTINGENCY PLANNED	SMALL	MEDIUM	LARGE
SMALL	695 (0)	391 (2)	158 (67)
MEDIUM	681 (2)	399 (0)	195 (35)
LARGE	.513 (35)	311 (28)	284 (0)

CHART 54

We call Chart 59 Contingency Planning A. Originally we planned to have a whole series of charts illustrating many different ways of doing Contingency

Planning but on second thought it seemed a little unreasonable to do so. The reason for this is that the problem of Contingency Planning is really too complex to be illustrated in our simple model. It usually turns out that the best that can be done is simply <u>to be aware</u> of the multiple situations possible and the multiple objectives in which one is interested. One then designs as good a system as possible and tests the performance of this system under all reasonable objectives and circumstances and even a few unreasonable ones. He then improves the design where it seems unsatisfactory, if necessary, at the expense of losing a little performance in situations where it is more than just satisfactory.

<u>A Better Contingency Design</u>

However, we should look at one more case. Actually we were a little unfair to the enemy in this last analysis. We gave him three fixed forces and designed against these fixed forces. The only freedom he had was to choose his best operating type. This is probably insufficient. After all, it probably takes longer to put up our defenses than for him to procure his offenses. Therefore, we should allow for the possibility that he may design his offense against our defense. It is probably best to give him a fixed offensive budget and let him allocate it in the best possible way for him. If we do this, Charts 55, 56, and 57 become Charts 60, 61, and 62, respectively.

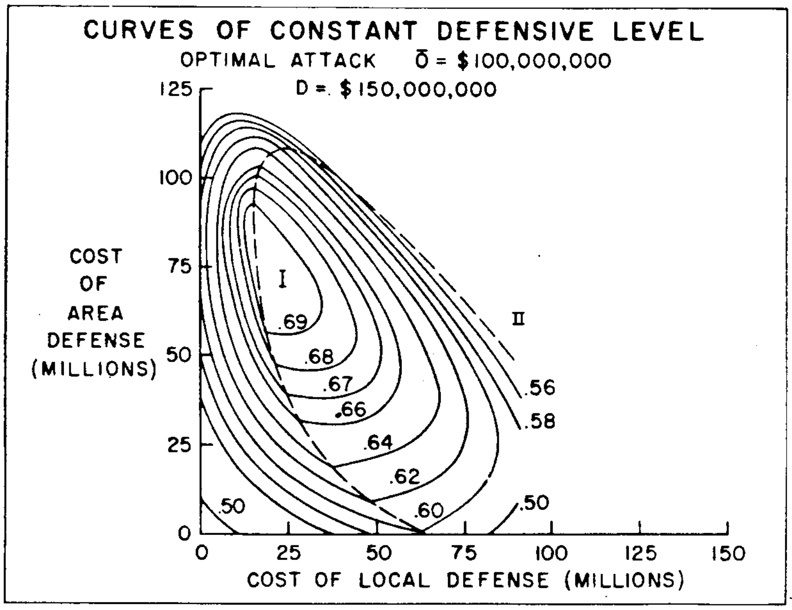

CHART 60

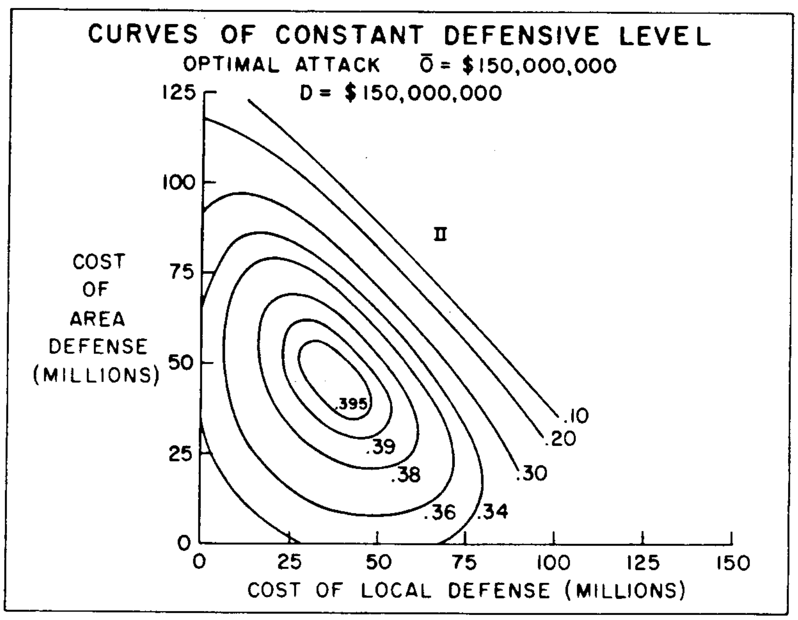

CHART 61

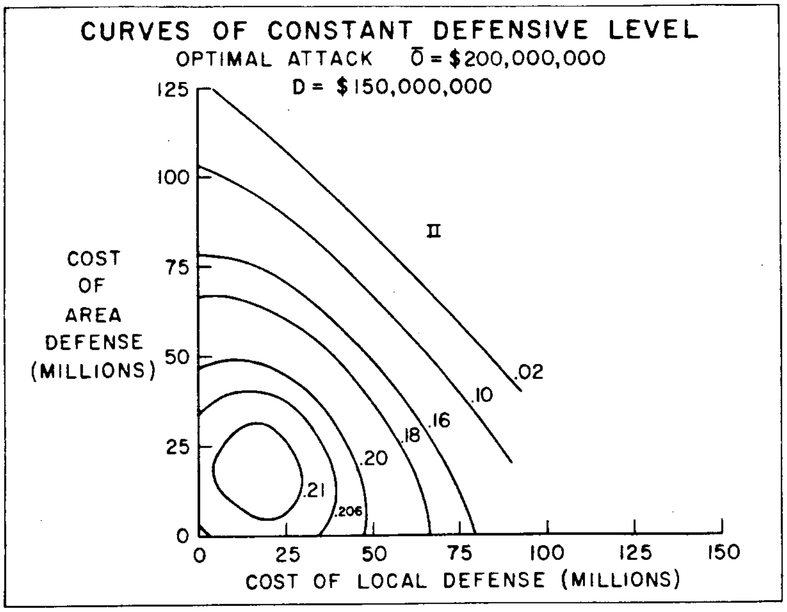

CHART 62

There are several things of importance to note. On the whole, Charts 60 and 61 are not very different from the old 55 and 56. The reason for this is that even though the old attack was not optimized, the specific allocation between bombs and planes is not very important (as we know from our studies in Part I). One can, however, overdo this philosophy and give the enemy a completely silly allocation. We did this in his large attack by giving him 24 planes and 16 bombs. Now, we already know that, when one is attacking with a lot of planes, one doesn't really need many escorts. The bombers themselves effectively saturate the active defense. In fact in this case the enemy has so oversaturated our active defense that we didn't bother to spend any money at all on it and put everything into passive defense. Under these circumstances his escorts are completely useless because he had no active defenses to circumvent. If he adopts a slightly more sensible policy and buys fewer planes and more bombs, we must have at least a small amount of area and local defense. If he adopted an even more sensible policy of buying more effective bombs and poorer planes, we would find it even more necessary to have active defenses. In other words, we can take some liberties with the allocation of the enemy's money, but liberty is no excuse for licentiousness. There is a difference between flirtation and seduction.

Soft Spots and Time Phasing

We have one last topic to discuss before finishing Chapter 3. Let us ask how our defense system will perform if the enemy happens to develop tactics outside our model[1]--for example, an air-ground missile. The air-ground missile is such that it can be launched from a point beyond our local defenses. The effect of this is that our local defenses no longer play a role. This can be serious. It means, in fact, that the money spent for local defenses has been thrown away. However, this new tactic is not entirely free to the enemy because his CEP (aiming error) goes up when he uses it.

We will assume that because of the increase in his CEP the value of our shelters goes up by some factor "m". The value of "m" will depend upon how far away the enemy is when he launches his air-ground missile. However, for the short range local defenses which we have been considering, we will take the factor "m" equal to 1.

Chart 63 shows us what the new tactic accomplishes. If we allocate our defense budget as the contingency planning study indicated, then our survival under the new attack goes down appreciably. In fact, even in the small attack this tactic means the difference between an almost certain retaliatory power remaining and a rather doubtful one. It is, therefore, very perilous to ignore

[1]At this point the reader may cry foul. He probably feels that we have no right to introduce an extraneous tactic. After all, we promised to stick to a definite model. We did, but we want to show that it may be disastrous when the enemy doesn't.

SENSITIVITY TO THE MODEL

D_A = \$ 30, 000, 000
D_L = \$ 30, 000, 000
D_S = \$ 90, 000, 000

	SURVIVAL	
ATTACK	CLASSIC ATTACK	NEW ATTACK
SMALL	.67	.48
MEDIUM	.39	.23
LARGE	.24	.16

CHART 63

the possibility that the enemy will develop tactics outside the model.

What is our counter to his tactics? One counter might be not to spend any money on local defense. Another might be to buy a longer range, more expensive local defense. Whatever counter we choose, we must keep in mind that the enemy is still always free to use the old tactic. The new tactics and their effects are shown on Chart 64. He must, even with the air-ground missile, still penetrate our area defenses, so they are just as effective as previously. This is symbolized by the fact that the money represented by area defense D_A is D_A. However, since he no longer comes within the range of our local defenses, the money spent on local defenses is worthless (at least as far as actually shooting him down is concerned). Lastly, the money spent on the shelters is actually more valuable. He is now delivering his bombs farther away and the accuracy of delivery is decreased; therefore, the effectiveness of his bombs is reduced.

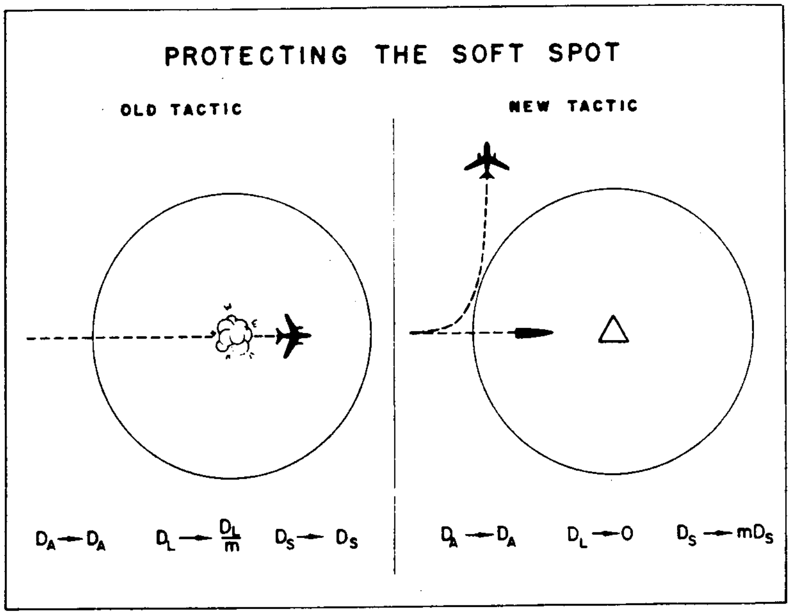

CHART 64

If the effectiveness of his bombs is reduced enough, the enemy may prefer to go ahead and drop his bombs in the old customary way. However this does not mean that we have not lost something. We have bought a more expensive local defense system, so we get fewer planes shot down per dollar. This is shown on the first part of Chart 64. The money spent on area defenses is still as good as before. The money spent on local defenses is less effective, because we have bought a longer range missile than we would need if we knew he was going to come in. We have assumed, for the sake of simplicity, that the factor by which the effectiveness of local defense money is reduced is the same "m" we used previously for increasing the value of shelters. Lastly, he no longer suffers a loss of accuracy, so the value of the money spent on shelters is the same.

It is important to notice that the choice of the factor "m" is up to us. We can try to develop (and spend money for) as long a range local defense missile as we wish. He has to stand off only as far as we force him to.

The last question is what should we do? Well, a lot depends on the pattern of information that we think exists. Let us assume that the enemy knows everything that we are doing. In this case, he is practically certain to choose what is best for him. Our alternative must be to design a defense system which is indifferent between the two attacks so that we don't care which he uses. This is the old problem of dividing a cake. If two children want to divide one in an even way, a good technique is to have one child do the cutting and the other choose the piece. The one doing the cutting does not make either piece too big because he knows that the other person would then choose it.

Chart 65 shows such an approach to this problem. (We have gone back to giving the enemy a fixed attack because it is really too complicated to let him optimize his attack, at least for this kind of discussion, and we decided that it was not important anyhow.) We first found a value of "m" for every budget allocation, which made us indifferent to whether or not he uses an air-ground missile. We then recalculated how well he does in

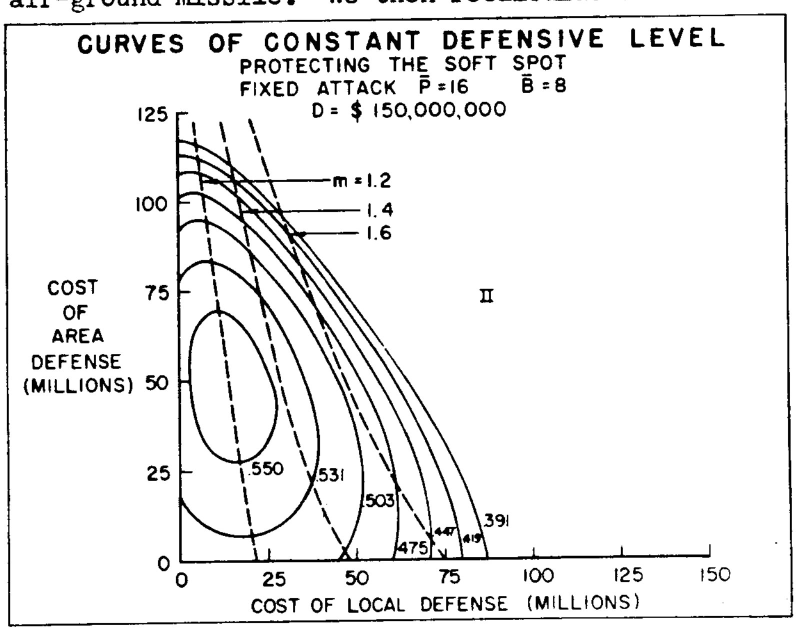

CHART 65

his attack and we also graphed a contour chart of his performance. On this contour chart we put the value of "m" needed for every allocation. We can now choose our allocation and value of "m".

Actually, we don't want to go into this in too much detail for two reasons. First, we have begun to push our model too far. As we extend the range of the local defenses, they begin to look more and more like area defenses. Our simple picture of forcing the enemy first through area defenses and then through local defenses breaks down. Second, one really should do this analysis with Contingency Planning, and we don't believe enough would be learned in this particular part of the example to justify the work.

One further thing should be mentioned. We have assumed that the enemy had an air-ground missile available. It may turn out that Intelligence tells us that he does not actually have one but is planning on developing it. If we are willing to rely on this information we are between the situation described in Chart 64 and the old situation. Presumably one would compromise by buying a more expensive local defense missile than one really needed right then, but not quite as expensive as one would want when the enemy actually has his air-ground missile. That is, we compromise a little, but not too much.

CHAPTER 4: THE TWO-SIDED WAR

Asymmetries

We should realize from the very outset that it is logically impossible to design a satisfactory system (except possibly in a mutual deterrence sense) based on a symmetric situation. For example, an air defense team will sometimes be given the responsibility of designing a system which, with High Confidence, will guarantee that the enemy will get less than 20% of his bombers through. In the same building, you will have an air offense team which is given the responsibility of designing a system that guarantees that at least 80% of our bombers get through his defenses, also with High Confidence. The astute reader will notice that even though the numbers 80 and 20 add up to 100, they are in this case completely incompatible, having a discrepancy of 60%. If our air offensive team attacks our air defensive team, there will be blood. At least one of them is going to be unhappy, probably both. The two problems cannot be simultaneously solved on the

SOME TYPICAL ASYMMETRIES

1. TECHNOLOGICAL SKILL
2. MILITARY SKILL
3. ALLIES
4. GEOGRAPHY
5. RESOURCES
6. INTELLIGENCE

CHART 65A

basis of technology alone. It can be done only if we can exploit some asymmetry between us and the enemy. Let us consider what some of these asymmetries might be, say, between the U.S. and the U.S.S.R.

1. Technological Skill

We can try to have better equipment than he has. On the whole, we think we are from one to five years ahead of the Russians. It is clear, however, that we cannot rely on this kind of superiority in all fields. Furthermore, we must really hop to keep ahead, if we are ahead. It cannot be done passively. However, where there is good evidence of a lead, it is probably a mistake not to exploit it. The thing works in reverse, too. We have to allow for it when there is a chance that he is ahead.

2. Military Skill

This is one of the things which, on the whole, Systems Analysts are forced to ignore, but on which policy makers often (and sometimes correctly) place much reliance. For an example of how important it can be, consider the Germans and the French in World War II. It seems to be true that the French had more and better tanks than the Germans, about as many planes, and, given their fortification systems, seemingly at least as good an army. The Germans were just better in the imponderables. Probably nobody could have predicted the German break-through and the completeness of their victory, unless he had given the Germans a great deal of credit for being better militarily. Of necessity, the analyst who is working for a High Confidence system should refuse to rely on the enemy's being incompetent militarily, unless there is very strong evidence to support such an assumption. Not only is it just too easy to indulge in wishful thinking but the Systems Analyst is looking to the future.

There is an important difference here between the current-policy maker and the planner. Conceivably the former, looking at the current situation, may wish to base his defense on the belief that the enemy is unskilled. The latter, making plans for systems with long lead times, simply cannot rely on the enemy remaining unskilled or stupid.

3. Allies

One of the important assets which we have is that, on the whole, the technically competent part of the non-Russian world is very closely associated with us. The Russians, of course, have the satellite nations on their side. Exactly what degree of dependence should be placed on allies is always a very touchy question. Presumably one tends not to rely on them for the High Confidence measures except in certain special cases; e. g., Canadian defense lines. One should, however, design his system to get advantages from their existence and conversely to contribute to their objectives.

4. Geography

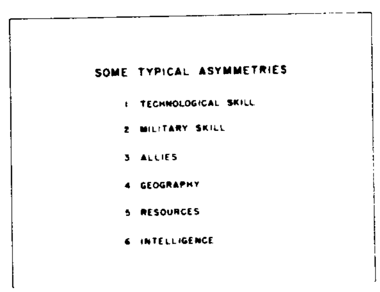

CHART 65A

One of the major asymmetries between us and the Russians is geography. The Russians have many millions of square miles more land than we have. This gives them certain advantages of depth. Against this must be balanced the fact that we have reasonable control of the sea, militarily valuable allies, overseas bases, distant early warning lines, etc.

5. Resources

We have a very much larger gross national product than the Russians, but they have a greater willingness to make the sacrifices necessary to keep a modern military mechanism in existence. They have not hesitated at any time to make huge expenditures when they felt it desirable. While we also have made large expenditures we seem to be much less willing to allocate money or manpower to military purposes. We are also much more swayed by the pressure of outside events. For example, in the postwar period our military budgets have been as low as 13 billion a year and as high as 50 billion. Most of the fluctuation came as a reaction to Russian moves.[1]

6. Intelligence

It is plainly easier for the Russians to get good intelligence about us than for us about them. This is probably true in all phases of intelligence activity. However, from the view point of High Confidence systems, it might easily be true that even the Russians could not rely on keeping things secure from us.

They might not, for example, go to large expenses to move their air bases around continually because the money would be wasted if they had only one properly placed defector.

The above list could be continued further, but we hope that the point is clear. Any kind of reasonable work in this field must include a great deal of attention to exploiting the asymmetries that are favorable to us and to ameliorating the ones which are not. However, in our study today, we shall not only be two-sided, we shall be symmetrical. We shall ignore what is in effect the heart of the problem and assume that the enemy and we have essentially the same capabilities and weaknesses, except that we shall give ourselves a slightly larger budget.

[1] There is more discussion in the Miscellaneous Comments chapter (Part Three) about the relative U.S., Russia, and NATO military resources.

Objectives--Two Kinds of Deterrence

When we discussed different kinds of objectives in Part II, we referred to the objective of achieving a 30% level of destruction as a deterrence objective, and we called a 60% level of destruction a "win-the-war" objective. We may have been somewhat misleading there. While it is true that the target system you pick is at least partly determined by whether you are trying to "deter" or "win the war", the real deterrence, at least against an enemy attack on you, is not what you can do if you have the first strike. It is, instead, your ability to survive the enemy's first strike and then strike back.[1] Chart 66 illustrates the situation.

TWO KINDS OF DETERRENCE
PERCENTAGE OF RED TARGETS DESTROYED

SYSTEM / WHO STRIKES FIRST	A	B	C	D
BLUE	.90	.40	.90	.90
RED	.20	.40	.40	.60

CHART 66

[1] Of course, if the enemy feels any doubts about his ability to surprise you; i.e., if he thinks that you might obtain strategic warning and forestall his attack, then the ability to strike first may also act as a deterrent. Also vice versa. For this reason, even Force A (to be discussed) may have some marginal value as a deterrent against his planning a surprise attack.

Assume, for example, that you are trying to choose between two forces, one which we will call Force A, and the other Force B. A is such that if you strike first, you will be able to get 90% of the enemy,[1] but if he strikes first, you will be able to get only 20% of him when you strike back. B is a different kind of force. It gets 40% of the enemy whether you or he strikes first. It is indifferent to how the war starts.

It is clear that if there is a tense situation and he is trying to decide whether or not to launch a surprise attack on you, then Force A will not deter him from launching such an attack. Quite the contrary, it will induce him to do it. He will argue that not only is there a tremendous payoff to him if he gets the first strike but that in fact he had better hurry up and get it. He feels that you must also realize the importance of going first, and if he does not move right away it may be too late. In fact, if the situation is as pictured, you could conceivably feel forced to take action yourself.

In other words, the A type of force might be called an "invitational force." The invitation is so strong that it practically commands you and the enemy to get in your blows early. Force A is a very unstable influence.

Force B, on the contrary, is a stabilizing force. The enemy looks at it and says, "I had better not start a war because if I do he will destroy 40% of me and that will hurt. Also, I get no advantage by moving first, because this thing will perform just as well regardless of who makes the first move. For the same reasons, he is in no hurry either, and we can both afford to sit and think about it."

Let us now look at a third force, Force C, also on the chart. At first sight it seems to combine the best virtues of Forces A and B. It gets 90% if

---

[1]We are being deliberately ambiguous by not specifying whether it is 90% of his industry, population, or military.

you go first, and 40% if he goes first. It is clear, however, that Force C is not as good a deterrent, at least in the sense in which we have been using the word, as Force B. Even though it does the same damage as Force B when the enemy strikes first, the enemy is still (relatively speaking) anxious to start the war. If he waits, you may start it and then he will really be badly off.

Therefore, if you are thinking only of deterring the enemy from a direct attack, you would probably prefer Force B to Force C. Furthermore, if you want to improve your offensive capability, it probably behooves you at the same time to improve your defensive ability, as illustrated by Force D. The extra strain on your deterrence is then matched by your extra strength.

Actually, of course, most people would argue that C is a better force than B. We believe that these people have a great deal to say on their side, though exactly how the thing comes out numerically, we are not sure. That is, once again we seem to be in the field of intuitive judgment rather than considered opinion.

The reason we might prefer C to B is that we are interested in doing more than just deterring a direct attack on ourselves. We are equally, or almost equally, interested in deterring the enemy from various kinds of provocative actions. For example, we don't want him to attack Iran or Korea or to make ultimatums against our friends in NATO or, in general, to misbehave. We shall call this kind of deterrence a Type II Deterrence as opposed to the previous kind which we shall label a Type I Deterrence.

Type II Deterrence is important. It is after all true that both the first and second world wars were started by the Allies' presenting the other side with ultimatums after the other side had put the Allies in an intolerable position. That is, in the first World War the British declared war on the Germans and not vice versa. Similarly, in the second World War the British and French declared war on the Germans after their invasion of Poland.

In the present very tight situation, we prefer not to correct that kind of misbehavior after it happens. We want to deter it from happening. To some extent this can be done by making the potential aggressor apprehensive that, if he pushes us too far, we just might take a drastic action of some sort. For example, we might try to convince him that there is at least a small chance that we are willing and able to declare war or that we are at least capable of putting ourselves into a position where we would be willing to declare war.[1]

---

[1]It is clear, of course, that many things restrain a potential aggressor. Among other effects, he may be afraid of:

- internal reactions
- losing friends or antagonizing neutrals
- creating or strengthening coalitions of enemies
- lowering the reaction threshold of potential opponents
- diplomatic or economic sanctions
- not going to heaven.

Our only point is that as far as a strategic air force contributes to Type II Deterrence it is mainly because the aggressor is afraid to provoke:

- an attack
- a higher or (from his point of view) more dangerous form of alert
- an expansion of our force
- direct military support for the threatened country.

If we wish to have our strategic air force contribute to Type II Deterrence then it must at least seem to be capable of taking one or more of the above actions. Usually the most convincing way to seem willing is to be willing.

Sometimes people also talk about what might be called a Type III Deterrence, to deter the enemy from relatively minor, but absolutely important, provocations (Peripheral Wars). It seems to the authors that in some real sense the job of the military forces here is as much actual ability physically to prevent or correct these provocations as they occur, as to deter them from happening. Insofar as there is deterrence, it is essentially the same kind as exercised by the Police Department against crime. Undoubtedly many people are deterred from committing a crime, but deterrence is not expected to be

(footnote continued on next page)

Now we should be aware that, if the enemy is afraid to attack Europe for fear we will in turn attack him and yet still wants to attack Europe, he might start his attack with a surprise blow at our own SAC. That is, insofar as we have any effective deterrence of the second type, we need even more deterrence of the first type.

There is an interesting side-light on the whole question of deterrence of the second type which should be mentioned. One may be perfectly willing to have certain kinds of powerful offensive forces, even if they are extremely vulnerable. There are at least three reasons for this.

100% successful. On the contrary, the police expect to go out and arrest and punish criminals in addition to deterring borderline cases from committing crimes.

Insofar as the enemy is deterred from these relatively minor provocations, it is probably true that the non-military things mentioned above are as important or more important than the kinds of things the military forces can provide. This is not to say that the military forces are unimportant in these cases, but simply to say again that their role is as much regulation and enforcement as deterrence.

We should also point out that insofar as we wish to rely on Type II or Type III Deterrence, even small and relatively ineffective forces from the military point of view can have some effect from the deterrence point of view. They make the enemy's aggressions more spectacular by preventing him from presenting us with a fait accompli that is psychologically hard for us to react against. Furthermore, if these relatively ineffective forces happen to be our own then we are almost automatically committed to react even if we would no longer want to. Both of these effects are sometimes called the "plate-glass effect."

1. As we mentioned, Type II Deterrence is at least partly measured by our strength when we go first. Therefore, even though these forces are vulnerable if the enemy strikes first, they may still make him wary of provoking us.
2. Even if the enemy believes that we will not go first, he cannot afford to rely on this belief. When it comes to long-term measures, he has to look at capabilities and not intentions. As a result, he may be forced to buy very expensive defensive systems in order to protect himself against our current or potential offensive forces. Any money that he is forced to invest in defense he cannot invest in offense. Therefore, our current or potential offensive forces can make a real contribution to our defense.
3. We may wish to use the vulnerable offensive force as a bargaining point. For example, consider President Eisenhower's "open skies proposal." It is not a disarmament proposal but a prelude to one. Its main purpose is to insure both sides against surprise attacks or trickery. Insofar as the Russians are not afraid of surprise attacks by us they may feel that it does not pay them, at least in the short run, to buy this proposal. In some circumstances, therefore, it might make sense to create forces whose main value would be their use as diplomatic pawns. We are willing to give them up for concessions from the potential enemy. In particular, by changing the deployment or alert status of special parts of the force one may be able to put real pressure on the enemy without making explicit ultimatums. This may be worth doing even if the special part is very vulnerable, as long as one's main forces are protected. (Of course, we may so aggravate our opponent by the creation of these

irritating forces that he is in no mood to bargain. Or what is sometimes more to the point, we may scare or antagonize our own citizens or allies.)

Let us, however, leave these realistic and somewhat unpleasant questions and go back to our toy. We will now consider the performance of different forces as measured by deterrence of the first kind. This means we will evaluate the effectiveness of Blue's force by how well he does after he has been attacked by Red. We will not look at Type II Deterrence because it would complicate our consideration too much. In any case, it is generally considered to be much less important than Type I in the world of today.

Both Sides Have Fixed Military Budgets

To begin with, let us assume that the total military budgets of both Blue and Red are fixed at some definite amount, $350,000,000 for Blue and $250,000,000 for Red. However, we will no longer assume given divisions of this budget for offense and defense purposes, but let each side allocate it in the best way possible for itself. Red will strike first and then Blue will strike back. We will measure success or failure for both sides by how much of Red survives Blue's retaliatory strike, in line with our statement that this measures Type I Deterrence.[1]

---

[1] It is only an approximation, and sometimes a misleading one, to evaluate one's success in defending an offensive force by the percentage of planes surviving. The important thing is to see how well the surviving force does when it counterattacks, taking full account of all the things that the enemy could do to prevent these surviving planes from attacking. Our analysis takes a step in this direction by evaluating the performance of Blue by how much destruction of Red he can do after Red's attack. However, our model is too simple to do this kind of evaluation well.

To repeat, the main problem is to decide what are reasonable allocations of each military budget between offense and defense. Blue's objective is to obtain a maximum destruction of the Red forces. Red's objective is to minimize the destruction of these same forces.

One way in which the best alternatives for Red and Blue might be explored is to start as we did in Part I and consider a reasonable number of possible alternatives for each side and then for each side to choose among these alternatives in such a way that its objectives are maximized. Charts 67 and 68 show a small but reasonable number of possible alternatives for this more complicated problem. We will assume, for the present, that the order of the decisions is as shown on the chart and that each participant is fully cognizant of all previous decisions.

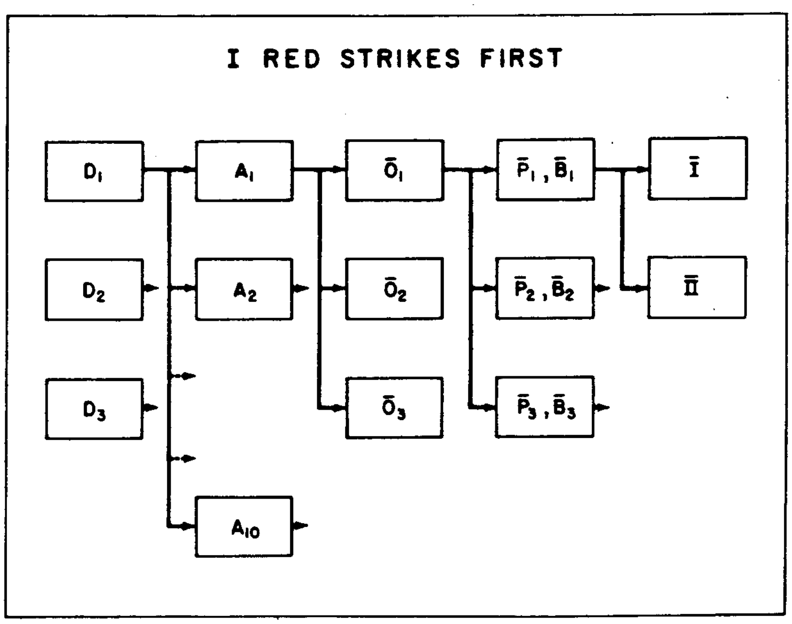

CHART 67

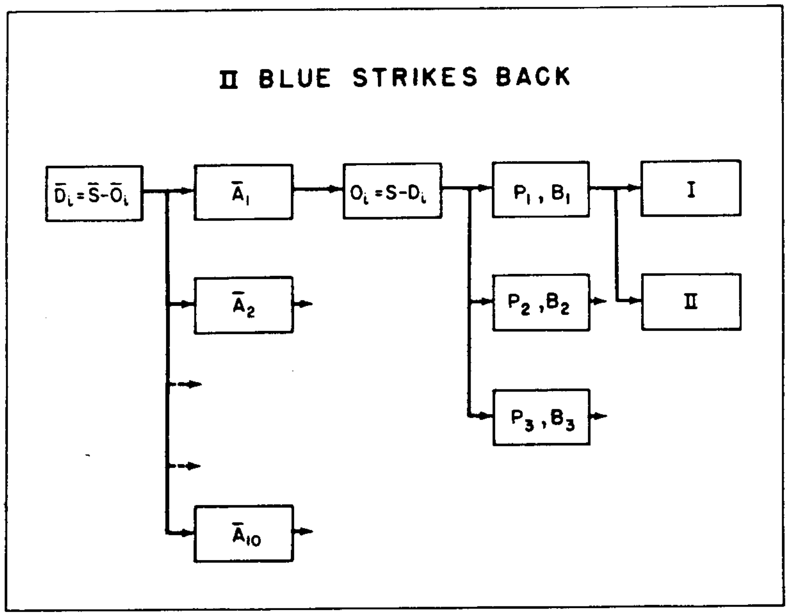

CHART 68

Chart 67 shows some of the different possible cases that might be considered in analyzing Red's strike. First Blue can choose among any of three defensive budgets. Then, for each of these defensive budgets, he can choose any of, say, ten allocations to area defense, local defense and shelters. The reason there are ten choices here rather than three is that there are now three different things to choose among, so it takes more cases to explore. Now Red can choose any of three offensive budgets. For each of the offensive budgets there are three allocations between planes and bombs. And for each of these, there is a choice of what tactic is to be used, Type I or Type II.

Let us now look at Chart 68, which illustrates Blue's retaliatory strike. Since Red has already chosen three offensive budgets, he is left with three defensive budgets. He can allocate them among area defense, local defense and shelters in ten different ways. Blue is left with three offensive budgets. He can allocate between planes and bombs in three different ways and then he can decide on one of the two tactics.

Now, even in this relatively simple situation, there are a large number of alternatives to be studied. In fact, multiplying them out, we have 3 x 10 x 3 x 3 x 2 x 10 x 3 x 2, which is 32,400 alternatives. Investigating each one of these and processing them is a major task. It is also a rather useless one because doing things in this way does not lend itself to performing the most important functions of Systems Analysis, such as taking account of uncertainty, doing contingency planning, considering the sensitivity to the model, and most important of all, training one's own intuition so that one can do a good design job. If these were not enough, there is the final consideration that in anything but a toy problem, this systematic approach would be quite impossible as the number of alternatives is just prohibitive. We have not included, for instance, the possibility of different military budgets, more alternatives at each stage, different orders in which the choices might be made, or different states of information.

The Approximations

However, we do want to get some feel for the two-sided situation. In order to do this, we must simplify the problem. We do this by using the lessons we have learned in the first three parts of the study and introduce the following approximations.

1. We will use an expected-value model.
2. Rather than study various allocations of Red's offensive budget, we shall assume that Red buys equal numbers of planes and bombs. In any reasonable situation this will not cut Red's performance by more than a few percent.
3. Since Blue is preparing for an attack by Red, he expects to lose some of his planes so he would not want to buy one bomb for every

plane. We shall assume that Blue buys half as many bombs as planes so that except in the most disastrous circumstances he will not have to leave any bombs home.

4. The defense allocations of both sides will be in the proportions that were determined in Chapter 3. We showed there that if these proportions were used the results were insensitive to the enemy's attack.

With the aid of these approximations we can now easily do the two-sided calculations and get some idea of how each side should allocate his budget between offense and defense. Once we have these allocations we can, if we want, (or need) redo the analysis of Chapters 1, 2, and 3, using these improved Offense and Defense budgets.

It is important to note that the approximations we made were not arbitrary, but have been shown to be <u>intuitively</u> reasonable by previous calculations. If we had not done these previous calculations we would not now be in a position to intuit which kinds of approximations are reasonable and which are not.

<u>The Two-Sided Calculation</u>

There is one further simplification we can make. We can draw the curves shown on Chart 69. These curves give the fraction of Blue destroyed when the ratio of the Red offensive budget to the Blue defensive budget is decided. In our simple model (and also in many complicated ones), the ratio of the two budgets determines everything independently of the actual budget levels. We have supposed that the defensive allocations are the ones determined in Part III, as we know this works well for a range of attacks. Red will, of course, use the tactic which appears best for him as shown

RM-1029-1 140

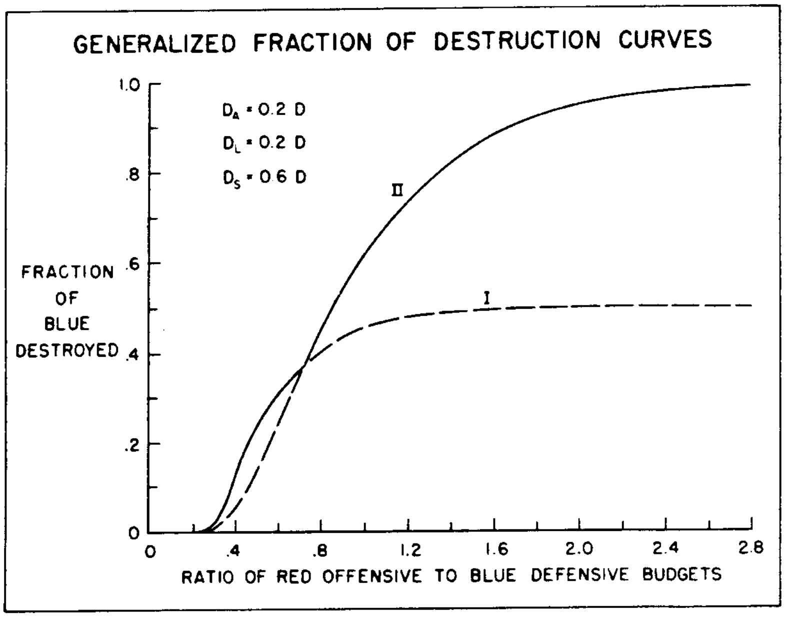

CHART 69

by the solid line on the chart.

This chart also can be used to determine the fraction of Red which is destroyed when Blue strikes back. Once Red has chosen his offensive budget, his defensive budget is then fixed, and likewise when Blue has chosen a defensive budget, his offensive budget is then fixed. However, part of the money spent on planes, and sometimes on bombs also,[1] will have been lost when Red strikes. If Red gets less than half of Blue's planes, then Blue will strike with all his bombs. If Red gets more than half of Blue's planes, then he will be unable to carry all the bombs he has bought.

Chart 70 shows a series of curves, one for each of six Red offensive budgets. Each of these curves shows the fraction of Red that will be destroyed, corresponding to the choice that Blue makes for his offensive budget. There are several things to be said about this chart. It is clear that if Red

---

[1]It is a typical weakness of many models that important components like bombs are invulnerable by assumption rather than by analysis.

RM-1127-1

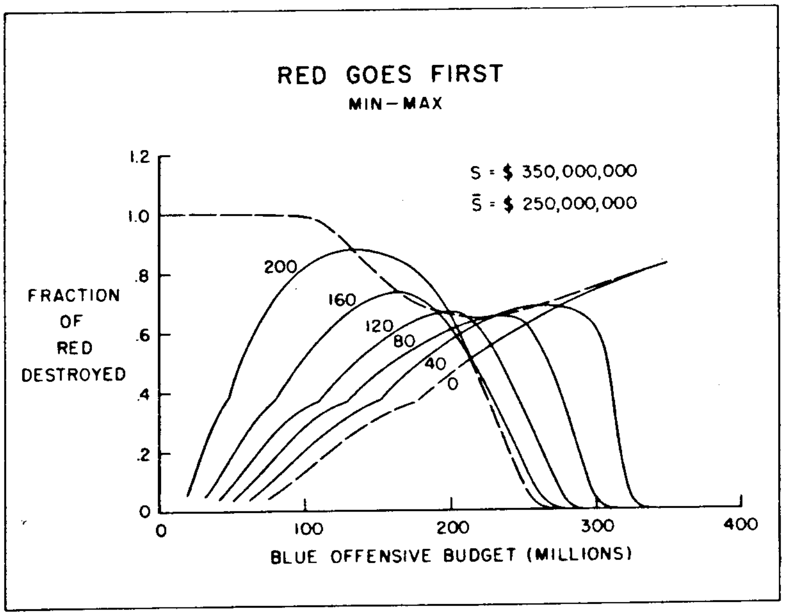

CHART 70

chooses and Blue knows what Red is doing, Blue can choose his offensive budget to obtain a value of destruction on the upper dashed line. That is, each time Red picks a budget he picks a curve; Blue should then pick the maximum point on this curve. Therefore, if Red is smart, he will pick the curve (around $90,000,000) which has the lowest maximum. This is the so-called min-max point (discussed further in the chapter on Game Theory). Blue will then make his offensive budget about $225,000,000 and destroy about 65% of Red's force.

The situation is different if Blue chooses first. For each Blue choice, the minimum destruction of Red is given by the curves comprising the lower dashed line. Red will now choose this minimum. Blue should choose, therefore, the maximum of the minima, the so-called max-min point. This would be an offensive budget of about $220,000,000. Red could then choose an offensive budget of either $0 or $120,000,000 and have to reckon with only 50% destruction of his force. The difference between this and

65% in the case where Red chooses first is considerable. But it is not as overwhelming as it would have been if Red in the first case did not choose the min-max point; or, if Blue in the second case, did not choose the max-min point.

In each of the two situations mentioned (Red chooses first, Blue chooses first), there are optimal strategies for both sides against an intelligent opponent. If either side operates far from that strategy, it can be very costly. The situation is qualitatively different when Red and Blue make their moves simultaneously. The discussion of this difference will be deferred to the Game Theory chapter. There we shall find that on the average the amount of Red's force destroyed in this new case will be between the outcomes in the two situations already discussed.

Chapter 11 on War Gaming (Part Two) also discusses methods that can be used to analyze the Two-Sided War.

By using suitable simplifications we have obtained some estimates of the outcome of the two-sided campaign. We have seen, as we expected, that the pattern of information, and the time sequence of the choices that are made, may be crucial. Also, we have seen that, if either side uses a poor strategy, it can be catastropic. However, even if the two sides act intelligently, the pattern of information still makes a difference, although not so much as before.

The important thing to notice is that questions such as these probably make a great deal more difference to the answers than the rather mild approximations that we made to simplify the analysis. As long as these important questions are handled by assumption, it probably does not make sense to work hard to eliminate the inaccuracy caused by the introduction of <u>relatively well understood</u> simplifications.

CHAPTER 5: EVALUATION AND CRITICISM

This finishes our sample study. The question immediately arises whether or not this whole method of working on things is really worthwhile. We believe it is easy to show that it is better than tossing a coin or crystal ball gazing, but this, of course, avoids the issue. The important thing is: Is it enough better to make any difference ?[1] We believe it is. Whether this is true or not has to be decided partly by current results and partly by a promise of things to come. It cannot be demonstrated by this kind of book.

The activity is still relatively new and many of the people involved are still learning their business. Our own organization has been in existence for about ten years. Over that time we have made some very substantial improvements in the quality of our work. While we could not state today that we can, in a routine way, guarantee a competent study on any subject, we do feel a certain amount of competency in a fairly large number of limited fields and also in at least a few large ones.

We should concede, though, that the question we claim to have been treating in this particular example, the question of the over-all force composition, is one of the most subtle and difficult ones. However, it is often relatively easy to find holes or gaps in a system and to make recommendations which one believes will fill these holes or gaps. If we are right, we shall then have an adequate or at least a better balanced system. It is much more difficult to make an over-all recommendation

---

[1]The unprofessional reader may wonder why we introduce this apologetic note; the professional reader will recognize its relevance, but will wonder why we brought it up.

We did it partly because we wished to be fair and indicate that there is of necessity such a question, and partly because we do believe that Systems Analysis is currently useful and that it will become more so in the future as people get more skilled. We even hope that by raising the question we increase slightly the probability that people will become more skilled.

for even a roughly "optimal" balance between different types of planes, bombs, missiles, etc. Therefore, the Analyst should usually look for adequacy and not optimality.

In any case, we believe that nothing in this book can be very definite or convincing about how well Systems Analysts can or cannot do this kind of work. Only a critical examination of past studies and a reasonable prognosis of future ones could do this.

Analysis Omits

We ought to make a few remarks, however, about some of the things that might be added to the sample study we have just finished to make it more realistic and, at least potentially, valuable. There is, of course, no format which one can write down which will guarantee a good study. There are,

ANALYSIS OMITS

1. RESEARCH? DEVELOPMENT? PROCUREMENT? OPERATIONS?
2. SCALE AND DETAIL
3. MIXED FORCES
4. UNCONVENTIONAL TACTICS
5. TIME PHASING
6. RE-EXAMINATION STAGE
7. UNSETTLED QUESTIONS
8. CONVINCING COMPARISON OF ALTERNATIVES

CHART 71

however, general remarks which one can make which will at least point out some of the items which might be considered. One does not, of course, include all the points we are going to mention in all studies. For some comments on the necessity to limit studies properly, see the discussion on Overambition in the Chapter on Ten Common Pitfalls (in Part Three).

1. Research? Development? Procurement? Operation?

The first question to ask is: what kind of recommendation is one trying to make; why are we doing the study? Occasionally a study is done solely for self-education. One simply wanted to know more about the subject. It is probably not right to think of such studies as Systems Analysis,[1] though they can be very valuable. Usually though, a study should be sharply affected by whether the question being treated is associated with Research, Development, Procurement or Operation.

Popular and semi-popular opinion to the contrary, the ability to do basic or semi-basic research in this country is not (in an absolute sense) in short supply. In addition, if we compare the cost of most kinds of basic or semi-basic research to other costs of the total defense system, we see that on the whole it is extraordinarily cheap. Ordinarily one can legitimately recommend doing it on very flimsy evidence. That is, one shouldn't have to show that the research being recommended is likely to be useful, but only that there is a chance that it may be useful.[2] Generally

---

[1]In our lexicon Systems Analysis means System Design. People who are interested only in analysis will have to find their own word.

[2]The alert reader will notice that we have left out the two important questions which are of most interest to people handing out funds:

1. What research should you do if you have a limited budget?
2. At what point would you recommend not increasing the amount of basic research?

(footnote continued on next page)

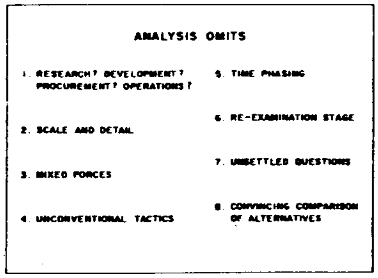

CHART 71

this kind of research should have only the loosest sort of guidance, except for monitoring to see that the people involved are professionally competent and that there are no glaring holes in the overall program.

Development is somewhat more expensive. But once again, one should not have to justify development program by proving that the item being developed should be procured. One simply tries to carry through, but in more detail, the kind of study that is done before recommending a basic research program. Here one might justify the program on the grounds that there are reasonable (in the case of expensive programs--not improbable) conditions under which this development could be useful. When the issue is crucial, we should only have to indicate that the nature of the problem is such that it is hard to show that the development won't be useful. In particular, we should do a great

---

On these points it should be mentioned that, in many cases, the problem of choice is not so difficult as one might think. It is clear that if one had a million development programs to choose among, and had no facts on which to base a choice until he had done some preliminary research on each program, then he might be in an undesirable position. In actual fact, it often occurs that there are only a few lines of development which are reasonable and that to pursue all of these lines to some reasonable point is quite cheap. There is then no problem of choice. One simply develops all the promising alternatives. In a study of a number of developments, B. Klein and W. Capron found several cases where the cost of the evaluation program that decided which development projects should be pushed was greater than the cost of pushing all of the likely looking projects.

For the hard and more usual case, where one has more uses for money than he has money, it is probably simply true that there is no general guide or general statement that can be made except the one that we have already made: The decision processes involved in Research, Development, and Procurement Programs are inherently sequential and it is very advantageous to exploit this sequentialness as much as is organizationally and politically possible.

deal of development simply to cut down lead time and provide insurance against uncertainty. One wants to be in a position so that if certain events (technical, political, or military) occur, we will be able either to guard against or exploit them. We should be willing to do this, even if the events are rather improbable, if they are important or if the cost of hedging is, relatively speaking, small. Much of the cost of development programs should be put down to giving this kind of flexibility.[1]

To quote a remark made earlier, "It may or may not be a mistake to develop something which is not procured, but it is always a mistake to procure something which is not developed."

Or referring again to the Time magazine article, it was only because the Navaho cruise missile was being pushed that we had the rocket engines for the ICBM. Actually, many people would have been willing to have a development program for big rocket engines which was not necessarily tied to any system at all. It just seemed very clear that large rocket engines are one of the commodities in which military systems of the future are likely to be interested.

In general, "state-of-the-art" development programs are good things. It turns out to be amazingly easy, ordinarily, to modify and wed already developed components.[2] A naive individual might think that if there were two independent programs which had not heard of each other and if at a later date one wanted to combine them into a single overall program, it might be impossible to do this. Sometimes it is, but more often it is not. The important thing is that if one insists on doing only complete system development programs, he will find himself

[1]One should probably differentiate between development slanted to answer performance and "state-of-the-art" questions and development slanted to making production prototypes. The former is much cheaper and should probably be treated more like basic research. Unfortunately there is some tendency to ignore this difference and think of all development programs as being slanted directly to developing production prototypes. This makes much development that should be thought of as answering performance and "state-of-the-art" questions more expensive and time consuming than necessary.

[2]We can hear the engineers swearing now! However, see note 1 above.

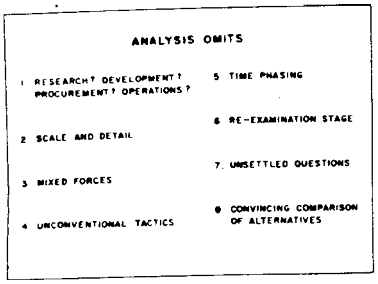

CHART 71

having a much smaller number of marriages of convenience to arrange. Historically, exactly this type of amalgamation often turns out to be important. Even if one knows that he wants a complete system he should be cautious in freezing the component design too early. There are a lot of errors made in predicting performance of components and one often achieves his goal faster by first doing bread-board work to answer state-of-the-art questions than by trying to do final design immediately.

The procurement problem is, of course, quite different. Procurement is expensive but has the virtue of involving much shorter time scales. When one has come to the step of being ready to procure something, the environment and context are relatively firm. There may be a lot of uncertainty left, but hopefully both the alternatives and the contingencies are reasonably limited. One should then be able to make just about the kind of study we have indicated in the example. Presumably the analysis should probably be pretty numerical all the way through, i.e., contain practically no qualitative arguments and conjectures except those that involve the kinds of fundamental questions that we necessarily relegate to the sphere of policy and intuitive judgment.

An operation study, of course, is almost completely concerned with detail, but there one has the saving grace of being restricted to current or immediately procurable hardware. For this reason, even though the amount of detail has gone up enormously, the number of alternatives has been correspondingly decreased. However, it is a typical weakness of current operational studies that they do not consider enough alternatives. For examples, consider some of the points which follow.

2. Scale and Detail

We talked about 20 planes, 2 airfields, etc. Actually, for some questions, one must talk about thousands of planes and hundreds of airfields. One may have to put in all the relevant asymmetries that were discussed in the early portion of Part IV including, for example, details of different capabilities, different targets, geography, etc. This is so obvious that it probably does not pay to dwell on it here. Unfortunately, it is often also not dwelt on where it should be. That is, not only must one be aware of the possibility that his model is inadequate to answer the questions he is asking but he must be aware of the need to make it adequate and not just to make some face saving apologies.

3. Mixed Forces

One of the most serious and common errors is that of arbitrarily concentrating the study on only one kind of bomb, one kind of plane, one kind of airfield, etc. Actually, of course, very few large organizations ever use only one type of equipment for all of their purposes. There is always a tendency to use specialized equipment for specialized problems. For the military too, it turns out that only by using a large number of different types of equipment and tactics can one really have the flexibility to meet effectively the many different objectives and circumstances. It is also very important either to consider a large variety of measures simultaneously or not to try to make one single measure handle all cases. While it is almost always necessary to limit the parts of the study that receive detailed treatment, it is desirable that this limiting be done intelligently and not arbitrarily.

4. Unconventional Tactics

There are two parts to this, the enemy's and your own. For the enemy, you are very interested in examining whether in plugging up one hole you

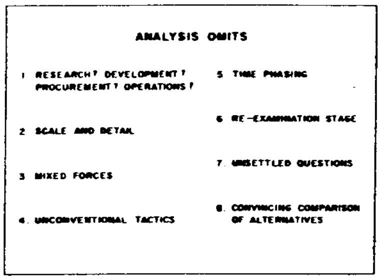

CHART 71

have forced him to another hole which may be equally disastrous to you. That is, you have to examine the recommended system and ask, "What can the enemy do to circumvent it?" In doing this examination one must give him all the freedom that physics, engineering and economics will allow.[1] All too often one finds studies which are designed against a specific enemy tactic rather than against the enemy himself.

There are strong psychological reasons for this. As long as a system has obvious holes, there is no reason for the enemy or us to consider subtle tactics. However, by eliminating these holes one has, in effect, forced the enemy to try to be clever. Under these circumstances he may consider tactics which once seemed far-fetched and improbable. Unfortunately, one may have to overcome a great deal of mental inertia (one's own as well as others) before one can take unaccustomed threats seriously, early enough to take effective action.

For our own side, as we mentioned before, the major objective of the Systems Analyst is not to analyze a given system but to design a system which

---

[1]We have deliberately left out social and political restraints. In the first pass at a problem one looks at capabilities and not intentions. While one may wish or need to modify his results to take account of such constraints on the enemy's behavior, it is often very hard to do this in a reasonable fashion. Usually, trying to exploit such constraints means dealing in Low Confidence measures. These of course can be useful. It is, however usually much more important to consider possible social and political restraints on one's own behavior than on the enemy's. Unless he explicitly and carefully considers such limitations, the analyst may find again that he is really dealing in Low Confidence measures even though from the viewpoint of technical capabilities it may look like a High Confidence measure.

will fulfill certain objectives satisfactorily. In doing this, he may also have to consider unorthodox or unconventional tactics in addition to recommending the development or procurement of new types of equipment.

5. Time Phasing

If it is the purpose of the study to recommend procurement and operation changes, it is often essential to take into account what one already has, and important to consider what one may want to have in the distant future. The recommended system should, therefore, consider the possibility of exploiting those things which one has and of having a salvage value for future systems. Also it sometimes makes a real difference exactly how the funds are to be disbursed on a year by year basis. While the Systems Analyst should not get into all the headaches of the executive, he should explicitly consider, if at all possible, everything which is important to his recommendations rather than make vague remarks about using or deferring to military or executive judgment.

6. Re-examination Stage

Most studies which have an impact start with a sort of re-examination stage where they investigate the current prejudices and concepts. It often turns out that a system designed for perfectly sound historical reasons is no longer effective because conditions have changed. However, the historical reasons for designing the system are embodied in the planning factors and tactics used in official studies. Therefore, the first job of the analyst is to get a reasonably accurate idea of what the facts really are. These cannot be obtained solely from official sources because it should be one of the major purposes of the study to review these numbers. It is also important for the study to generate on its own the numbers which

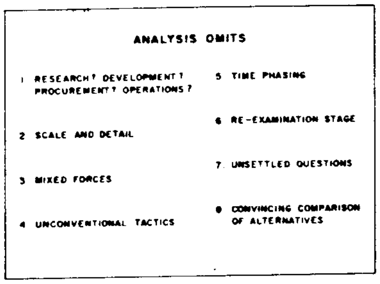

CHART 71

previously have been ignored or treated in an offhand way.

7. Unsettled Questions

It is not a sign of weakness but of strength to hold certain conclusions tentatively, particularly if one has indicated a program designed to settle them (insofar as they can be settled). It is important, of course, to take as firm a position as can be justified in a reasonable way and not to overemphasize the uncertainties. (One of the most common excuses for doing nothing is to say that nothing can be done until more information has been obtained. Sometimes the excuse-maker adds insult to injury by acting as a roadblock to getting more information.) But if a question is really open, then one should say so. While there is always an obligation to set up a program which will answer open questions or indicate that they are unanswerable, there is no obligation for the Systems Analyst to have a policy position in advance of the investigation. After all, he is not an executive.

In addition, most studies have a continuing existence and it is often wise to consider their interaction with the past and the future and to leave room for future development. Some of the recommendations that are made may well be made for the sake of future studies.

8. Convincing Comparison of Alternatives

Most important of all, a study that attempts to influence policy should have a convincing comparison of all the relevant alternatives. The kind of curves that we drew in our study are sometimes not directly to the point. Usually either with the aid of such curves or by some simple-minded considerations one succeeds in designing what he thinks is a reasonable system. The task is then to show (if possible) that, under any reasonable assumptions, the system designed by the analyst is to be preferred to any alternative systems which are being considered.

To give an example, let us assume that the analyst has come out for the mixture of area and local defenses and sheltered airfields that we arrived at ($30,000,000 for area defense, $30,000,000 for local defense and $90,000,000 for airfields). Let us assume that we already possess a $30,000,000 area defense, a $30,000,000 local defense and $30,000,000 in sheltered airfields. The analyst is then proposing to add another $60,000,000 to the sheltered airfields. We can call this proposal System A. There are, however, other proposals for spending the same money. The other proposals for systems might be as follows:

System B proposes to spend the whole $60,000,000 for local defense.

C proposes to spread the money equally between local and area defenses.

D proposes to spend the whole $60,000,000 for area defense.

Now it is the analyst's job, if possible, not only to design a good system but also to test his system in such a thorough way that he will be able to convince the proponents of the other systems that his system is better, or at least almost as good, as theirs.

It often happens in practice that people think they disagree on recommendations because they disagree on details of performance, when, in fact, one of the contenders could accept the other contender's assumptions on performance and still prove his case; that is, the analyst can use an *a fortiori* argument. It is essential for the Systems Analyst to see if he can do this.

Therefore, when we are comparing A with B, we shall assume that the technical performance of local defense is, in fact, as good and that of

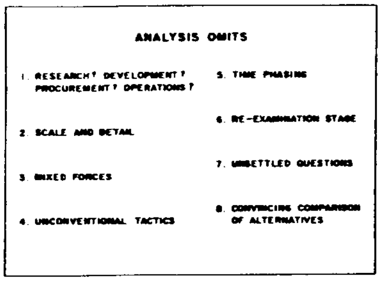

CHART 71

shelters as bad as the local defense enthusiasts think they are, and still see if A is better than B. Conversely, when we compare A and D, we shall assume that the area defense performs well and that the shelters perform badly. (If a pure $150,000,000 shelter system were also being considered then we would give the shelters a premium performance and the active defense a black mark.) In other words, in making preliminary expository comparisons, we bend over backwards to hurt our system and to help the alternative system. If it turns out after we have done this that we can still say we prefer our system then we are in a superb position to make recommendations.

One reason the above program can so often be carried through successfully is that many of the most successful and exciting analyses have about them a large element of "the Emperor has no clothes."[1] In such cases it is not surprising that one can go to extraordinary lengths in accepting the assumptions of one's opponents and still prove his point. In other cases people may think of the problem as being a choice between two alternatives when a clever analyst may be able to make recommendations that in effect put him in the position of being able to "eat his cake and have it too." Possibly because our moralistic culture tends to overlook and deny this possibility an ingenious analyst can often discover exciting and brilliant but "obvious" things.

[1]For those who have had no childhood: the phrase refers to Andersen's story, "The Emperor's New Clothes."

Sometimes we cannot go all the way in satisfying the B, C, or D enthusiasts. If we concede their presumably exaggerated performance claims for the gadgets they like and also concede their presumably undue pessimism for the gadgets we recommend, then in fact one or more of their systems may look better. Under these circumstances we have two alternatives. We can see how far we can go along with the opposition and conduct a so-called "break-even" analysis. In such an analysis we find just what assumptions we have to make about the important values in order to make the performance of the two systems the same. We can then simply ask people to judge whether these assumptions are unduly optimistic or pessimistic. Or we can try to make a more convincing case on what are reasonable assumptions. The best approach is generally to use both of the above measures.

We shall close this section by a numerical illustration of the above remarks. In Chart 72 we indicate how, using our assumptions, the four

OUR COMPARISON OF SYSTEMS

FRACTION SURVIVING ENEMY'S ATTACK

ATTACK	A	B	C	D
SMALL	.66	.58	.65	.68
MEDIUM	.38	.28	.34	.32
LARGE	.24	.07	.09	.08

CHART 72

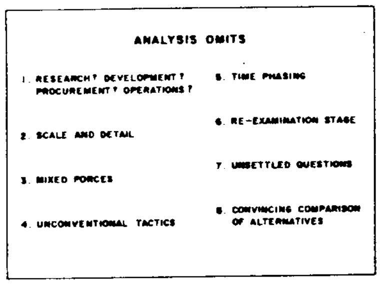

CHART 71

systems would compare under the three attacks. (In this and the next calculations we always let the enemy use his optimal tactic.) We notice immediately that System B looks very bad, System C is not so bad except for the heavy attack, and System D is superior in the case of the light attack, although it falls down badly in the case of the heavy attack. It is probably true that anybody who studies this chart (and believes it) would conclude that System A is superior. The trouble is we have to try to convince at least three different groups of protagonists, and these other people do not compute the same way we do, or even as each other does.

Let us assume that B thinks that the local defenses will be 10% more effective than we do (i.e. each $1.00 looks like $1.10 on our curves). C assumes both the area and local defenses are 10% more effective than we do, and D assumes that the area defenses are 10% more effective. They all also assume that shelters are 10% less effective than we have taken them to be (each $1.00 looks like $.90). B, C, and D also claim that their recommendations are different from ours because their assumptions about component performance are different. Let us see if this last conclusion is valid.

Chart 73 shows how the comparison comes out as calculated by each of these individuals. It is clear that even with his assumptions B has made a mistake. No matter what attack he is designing against he should not prefer his system to ours. C also has made a mistake since using his assumptions he does about the same as we do in the two light attacks and he does not have much insurance against the heavy attack. He really cannot claim any margin

INDIVIDUAL COMPARISON OF SYSTEMS

ATTACK	B'S CALC		C'S CALC		D'S CALC	
	B	A	C	A	D	A
SMALL	.58	.65	.67	.66	.70	.65
MEDIUM	.28	.36	.39	.37	.34	.36
LARGE	.07	.21	.09	.22	.07	.21

CHART 73

of superiority. D is in a slightly better position. His superiority in the case of the light attack is appreciable (but not overwhelming). Therefore, if he happens to think that the medium and heavy attacks are very improbable, so improbable as not to be worth considering, he may stick to his guns in preferring his system. It is important to realize that in conducting this argument, we don't have to persuade D that the heavy attack is probable; we only have to persuade him that he should not <u>rely</u> on it not occurring. That is, the only point we have to make is that even if the heavy attack is improbable, it is worth losing 7% on performance in the case of the light attack just in case this improbable but still devastating heavy attack occurs. <u>It is perfectly legitimate to talk this way even if one thinks the heavy attack is probable</u>. If we cannot do this, and if we also cannot convince D that his performance estimates are wrong, then we just cannot convince him. Possibly we ourselves are wrong.

Because there isn't much difference between us and D in the light attack, there is not much point in our conducting a break-even analysis. However,

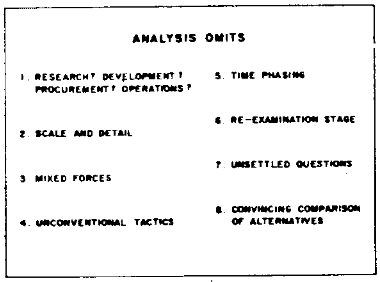

CHART 71

D may want to do one to find out what assumptions he has to make in order to look reasonably good in the heavy attack. His break-even analysis is shown in Chart 74, which shows how the comparisons are affected by using different performance assumptions.

The break-even analysis is really given by the single curve where k = 1. But we have

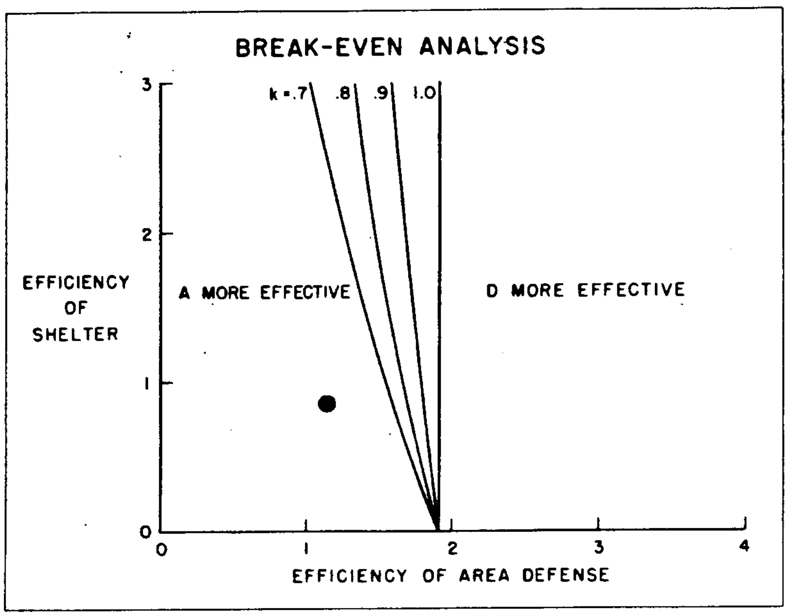

CHART 74

drawn other curves where k is the relative effectiveness of D and A under the heavy attack. These curves show D the assumptions he must make, even if he concedes that he will not do quite as well as A. The black dot shows where D's current assumptions are. (Note: It is just an accident of our model that the curve k = 1 is vertical.)

More than any other single thing, the skilled use of a-fortiori and break-even analyses separate the professionals from the amateurs. A good analyst should clearly separate the parameters (assumptions) into two parts; those to which he can afford to give quite pessimistic values and those to which he has to give "reasonable" or breakeven values. The analyst can then point out that *all one needs to believe in order to accept his recommendations is these few crucial assumptions*. If the audience accepts the assumed values of these crucial parameters as being reasonable or at least in the break-even region, then they *must* accept the recommendations. If the whole briefing is built around this idea, it is very surprising how even extremely unpalatable conclusions can be brought home.

In order to carry through the above program most analyses should (conceptually) be done in two stages: a first stage to find out what one wants to recommend, and a second stage that generates the kind of information that makes the recommendations convincing even to a hostile and disbelieving, but intelligent audience.

APPENDIX

12 Summary Charts

(from Chapter 2)

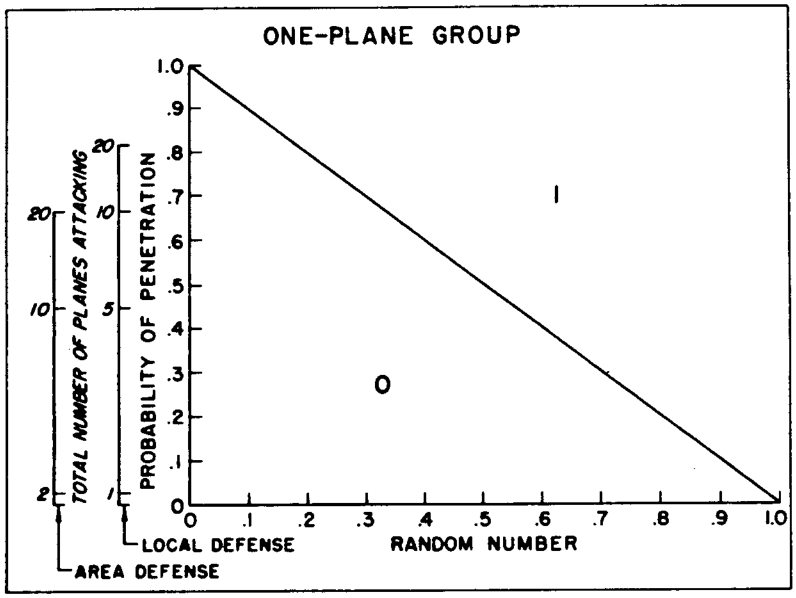

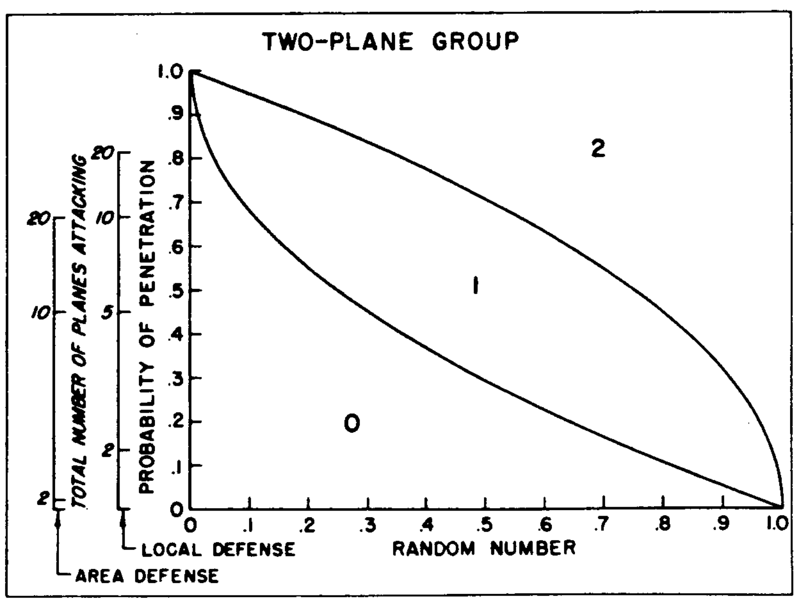

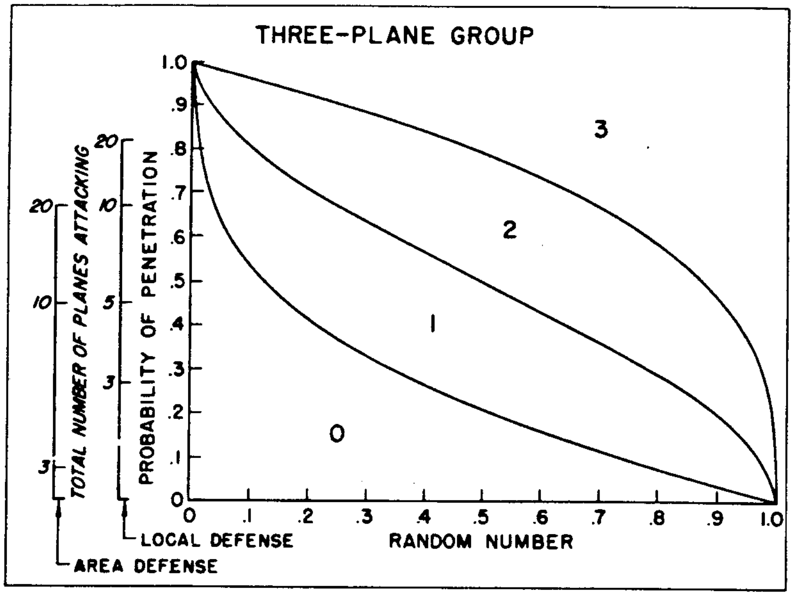
THREE-PLANE GROUP
PROBABILITY OF PENETRATION
TOTAL NUMBER OF PLANES ATTACKING
LOCAL DEFENSE
AREA DEFENSE
RANDOM NUMBER
0
1
2
3

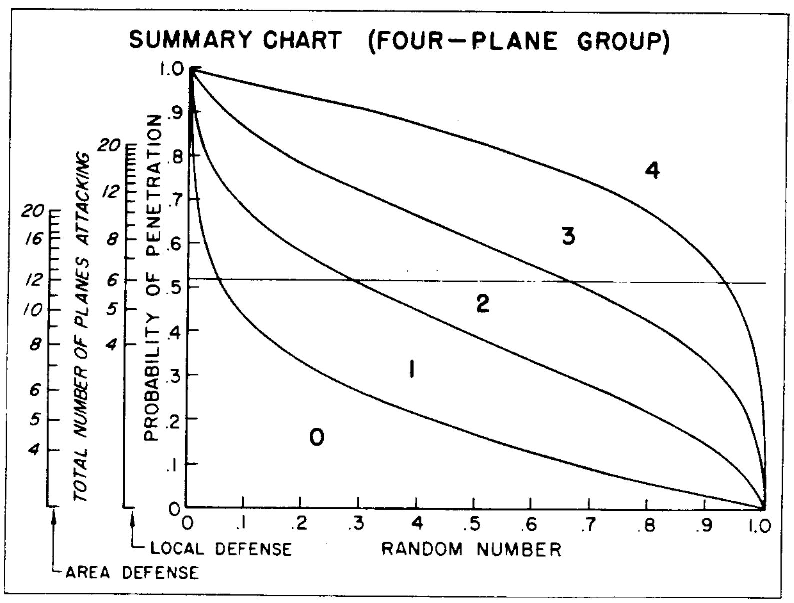
SUMMARY CHART (FOUR-PLANE GROUP)
PROBABILITY OF PENETRATION
TOTAL NUMBER OF PLANES ATTACKING
LOCAL DEFENSE
AREA DEFENSE
RANDOM NUMBER
0
1
2
3
4

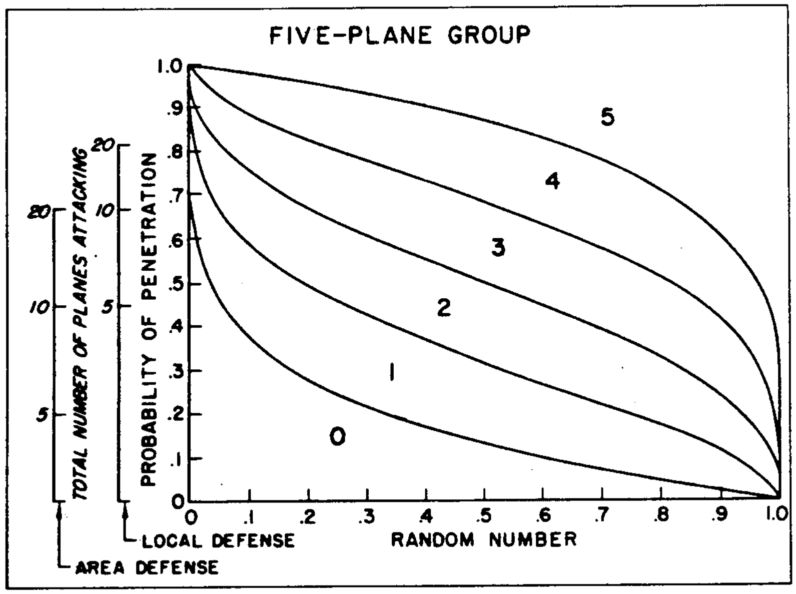

FIVE-PLANE GROUP
PROBABILITY OF PENETRATION
TOTAL NUMBER OF PLANES ATTACKING
LOCAL DEFENSE
AREA DEFENSE
RANDOM NUMBER
0
1
2
3
4
5

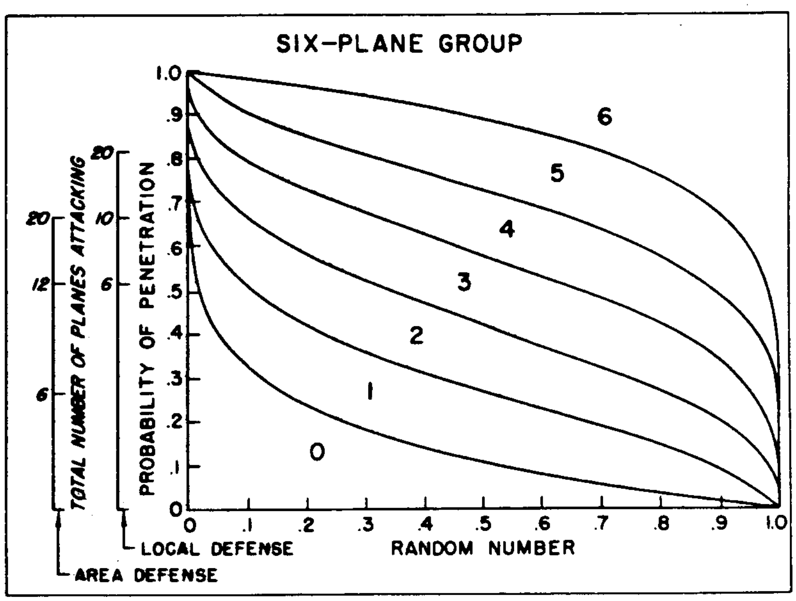

SIX-PLANE GROUP
PROBABILITY OF PENETRATION
TOTAL NUMBER OF PLANES ATTACKING
LOCAL DEFENSE
AREA DEFENSE
RANDOM NUMBER
0
1
2
3
4
5
6

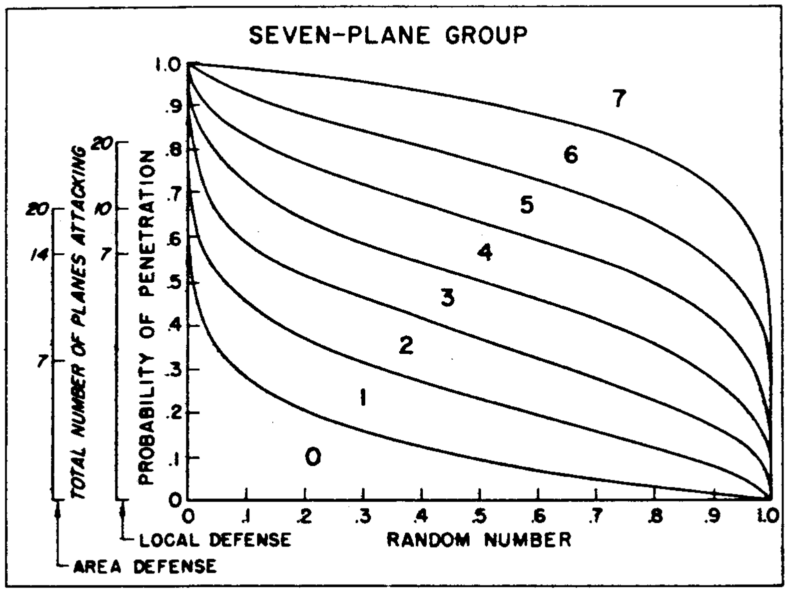
SEVEN-PLANE GROUP
PROBABILITY OF PENETRATION
TOTAL NUMBER OF PLANES ATTACKING
RANDOM NUMBER
LOCAL DEFENSE
AREA DEFENSE
1.0 .9 .8 .7 .6 .5 .4 .3 .2 .1 0
0 .1 .2 .3 .4 .5 .6 .7 .8 .9 1.0
20 10 7
20 14 7
0 1 2 3 4 5 6 7

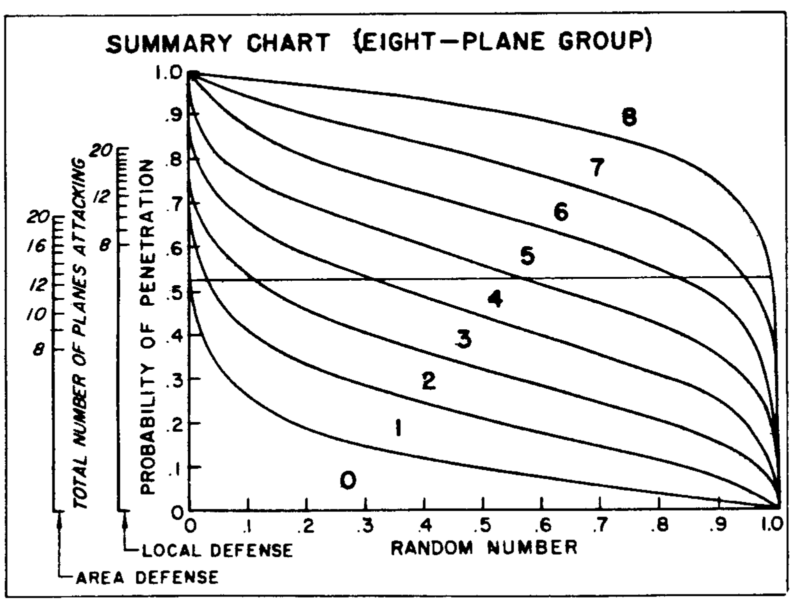
SUMMARY CHART (EIGHT-PLANE GROUP)
PROBABILITY OF PENETRATION
TOTAL NUMBER OF PLANES ATTACKING
RANDOM NUMBER
LOCAL DEFENSE
AREA DEFENSE
1.0 .9 .8 .7 .6 .5 .4 .3 .2 .1 0
0 .1 .2 .3 .4 .5 .6 .7 .8 .9 1.0
20 12 8
20 16 12 10 8
0 1 2 3 4 5 6 7 8

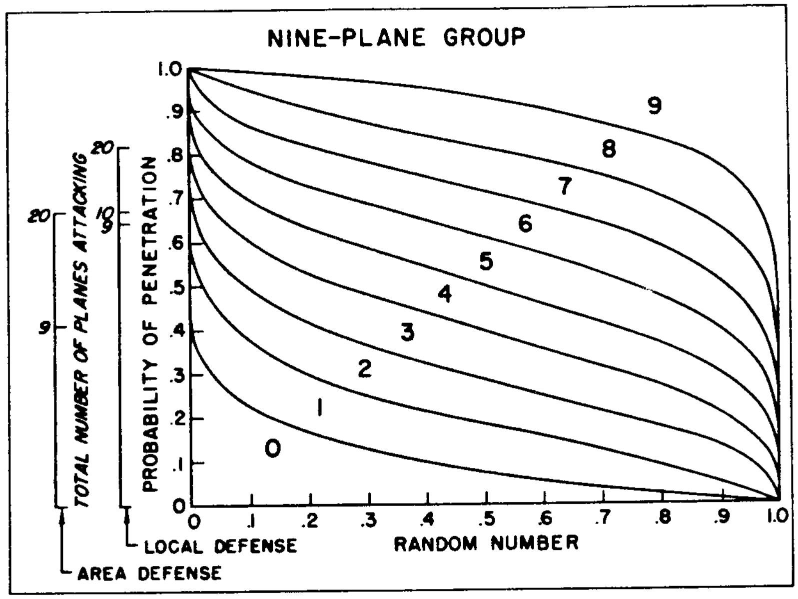
NINE-PLANE GROUP
PROBABILITY OF PENETRATION
TOTAL NUMBER OF PLANES ATTACKING
RANDOM NUMBER
LOCAL DEFENSE
AREA DEFENSE
1.0
.9
.8
.7
.6
.5
.4
.3
.2
.1
0
.1
.2
.3
.4
.5
.6
.7
.8
.9
1.0
20
10
9
20
9
0
1
2
3
4
5
6
7
8
9

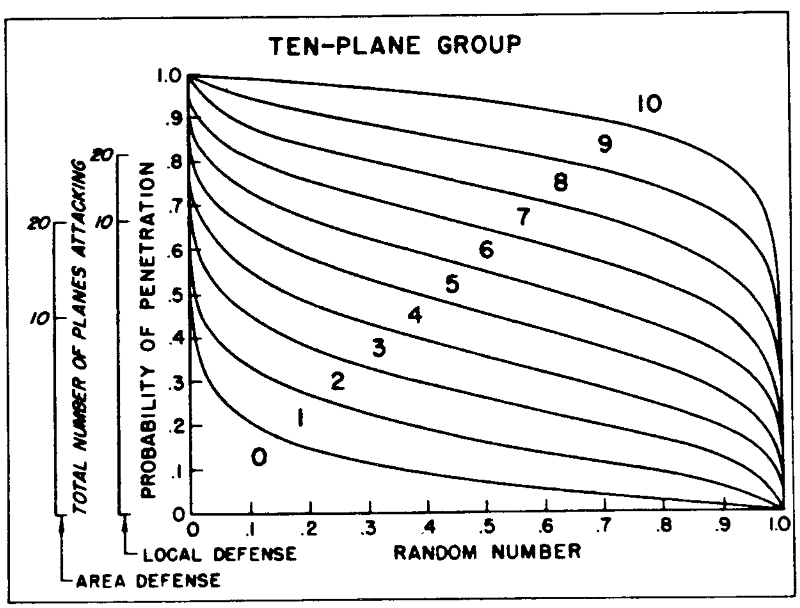
TEN-PLANE GROUP
PROBABILITY OF PENETRATION
TOTAL NUMBER OF PLANES ATTACKING
RANDOM NUMBER
LOCAL DEFENSE
AREA DEFENSE
1.0
.9
.8
.7
.6
.5
.4
.3
.2
.1
0
.1
.2
.3
.4
.5
.6
.7
.8
.9
1.0
20
10
20
10
0
1
2
3
4
5
6
7
8
9
10

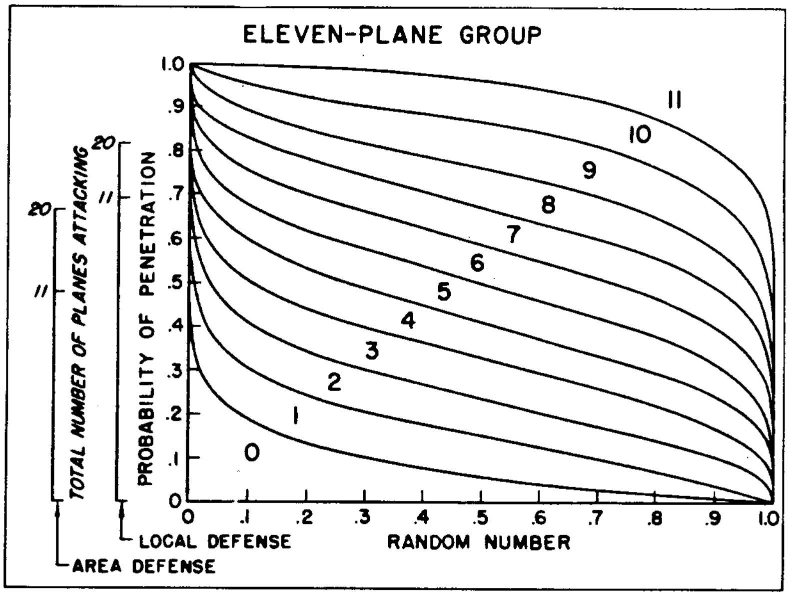
ELEVEN-PLANE GROUP
PROBABILITY OF PENETRATION
TOTAL NUMBER OF PLANES ATTACKING
RANDOM NUMBER
LOCAL DEFENSE
AREA DEFENSE
0
1
2
3
4
5
6
7
8
9
10
11
20
1.0
.9
.8
.7
.6
.5
.4
.3
.2
.1
0

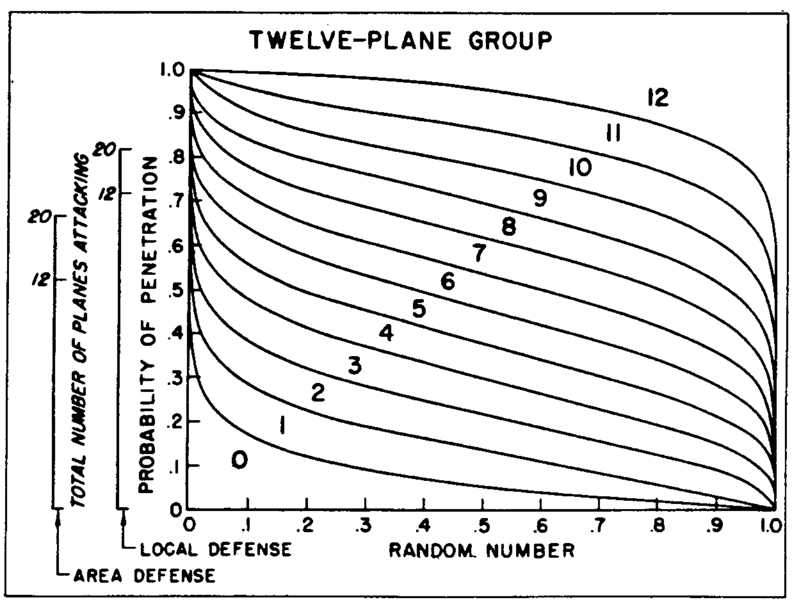
TWELVE-PLANE GROUP
PROBABILITY OF PENETRATION
TOTAL NUMBER OF PLANES ATTACKING
RANDOM NUMBER
LOCAL DEFENSE
AREA DEFENSE
0
1
2
3
4
5
6
7
8
9
10
11
12
20
1.0
.9
.8
.7
.6
.5
.4
.3
.2
.1
0

APPENDIX 2

The Model

There are five elements of the model:

1. planes
2. bombs
3. area defense
4. local defense
5. sheltered airfields

Planes and bombs are given costs of $5,000,000 apiece. A plane can carry one bomb, or can act as an escort.

The performance of the defenses depends on the budget given to each. Let

D_A = budget for area defense (millions)

D_L = budget for local defense (millions)

D_S = budget for passive defense (shelters) (millions)

P = number of planes in cell attacking area defense

P' = " " " " " " local "

p_A = probability of each plane penetrating area defense

p_L = " " " " " local "

b = number of bombs dropped on target

s = fraction of target surviving

The performance of the planes is given by:[1]

$$p_A = e^{-\frac{.16\ D_A}{P}} \qquad \text{(Chart 5)}$$

$$p_L = e^{-\frac{.08\ D_L}{P'}} \qquad \text{(Chart 7)}$$

and the performance of bombs by:[2]

$$s = e^{-\frac{b}{.04D_s}} \qquad \text{(Chart 8)}$$

These functions are graphed in Chapter 1 with $D_A = D_L = D_s = 50$.

---

[1]This formula can be seen to be intuitively reasonable as follows: Suppose F fighters are engaged in an air battle with B bombers. Under the assumptions

1. each fighter makes one pass at a bomber
2. the probability of kill is 1
3. the fighters choose the bombers independently of one another,

the probability that a bomber survives is

$$p = (1 - \frac{1}{B})^F$$
$$\approx e^{-F/B}$$

F is called the "kill potential" of the defense barrier, as the expected number of kills, for large B, approaches F. There are other assumptions leading to the same approximation, but these are often used.

[2]If the probability a target survives one bomb is 1 - p, then b independent bombs give a survival probability of

$$(1 - p)^b \approx e^{-\alpha b} .$$

This is the expected survival of the target. It can also be interpreted as giving the expected damage due to each successive bomb, as we have done in Chart 8.

APPENDIX 3

The Essence of Systems Analysis

(A Review of the Major Topics of Part One)

An item of equipment cannot be fully analyzed in isolation; frequently, its interaction with the entire environment, including other equipment, has to be considered. The art of systems analysis is born of this fact; systems demand analysis as systems.

Systems are analyzed with the intention of evaluating, improving, and comparing them with other systems. In the early days many people naively thought this meant picking a single definite quantitative measure of effectiveness, finding a best set of assumptions and then using modern mathematics and high speed computers to carry out the computations. Often (as discussed in Chapter 12, Ten Common Pitfalls), their professional bias led them to believe that the central issues revolved around what kind of mathematics to use and how to use the computer.

With some exceptions, the early picture was illusory. First, there is the trivial point that even modern techniques are not usually powerful enough to treat even simple practical problems without great simplification and idealization. The ability and knowledge necessary to do this simplification and idealization is not always standard equipment of scientists and mathematicians or even of their practical military collaborators. Some of the pitfalls that have been common in the past are discussed in Chapter 12.

Much more important, the concept of a simple optimizing calculation ignores the central role of uncertainty. The uncertainty arises not only because we do not actually know what we have (much less what the enemy has) or what is going to happen, but also because we cannot agree on what we are trying to do.

In practice, three kinds of uncertainty can be distinguished.

1. Statistical Uncertainty
2. Real Uncertainty
3. Uncertainty About the Enemy's Actions.

We will mention each of these uncertainties in turn.

(1) Statistical Uncertainty. This is the kind of uncertainty that pertains to fluctuation phenomena and random variables. It is the uncertainty associated with "honest" gambling devices. There are almost no conceptual difficulties in treating it--it merely makes the problems computationally more complicated. We discussed some aspects of this uncertainty in Chapter 2.

(2) Real Uncertainty. This is the uncertainty that arises from the fact that people believe different assumptions, have different tastes (and therefore objectives), and are (more often than not) ignorant. It has been argued by scholars that any single individual can, perhaps, treat this uncertainty as being identical to the statistical uncertainty mentioned above, but it is in general impossible for a group to do this in any satisfactory way.[1] For example it is possible for individuals to assign subjectively evaluated numbers to such things as the probability of war or the probability of success of a research program, but there is typically no way of getting a useful consensus on these numbers. Usually, the best that can be done is to set limits between which most reasonable people agree the probabilities lie.

---

[1] The Foundations of Statistics, by L. J. Savage; Social Choice and Individual Values, by K. J. Arrow.

The fact that people have different objectives has almost the same conceptual effect on the design of a socially satisfactory system as the disagreement about empirical assumptions. People value differently, for example, deterring a war as opposed to winning it or alleviating its consequences, if deterrence fails; they ascribe different values to human lives (some even differentiate between different categories of human lives, such as civilian and military, or friendly, neutral, and enemy), future preparedness vs. present, preparedness vs. current standard of living, aggressive vs. defensive policies, etc. Our category, "Real Uncertainty," covers differences in objectives as well as differences in assumptions.

The treatment of real uncertainty is somewhat controversial, but we believe actually fairly well understood practically. It is handled mainly by what we call Contingency Design and is discussed in Chapter 3.

(3) Uncertainty Due to Enemy Reaction. This uncertainty is a curious and baffling mixture of statistical and real uncertainty, complicated by the fact t[illegible] we are playing a non-zero sum game.[1] It is often very difficult to treat satisfactorily. A reasonable guiding principle seems to be (at least for a rich country), to compromise designs so as to be prepared for the possibility that the enemy is bright, knowledgeable, and malevolent, and yet be able to exploit the situation if the enemy fails in any of these qualities.

---

[1]As explained in Chapter 10, the terminology "non-zero sum game," refers to any conflict situation where there are gains to be achieved if the contenders cooperate. Among other things, this introduces the possibilities of implicit or explicit bargaining between the two contenders. The whole concept of deterrence comes out of the notion that the game we are playing with Russia is non-zero sum. (See pages to .)

To be specific:

- Assuming that the enemy is bright means giving him the freedom (for the purpose of analysis) to use the resources he has in the way that is best for him, even if you don't think he is smart enough to do so.

- Assuming that he is knowledgeable means giving the enemy credit for knowing your weaknesses if he could have found out about them by using reasonable effort. You should be willing to do this even though you yourself have just learned about these weaknesses.

- Assuming that the enemy is malevolent means that you will at least look at the case where the enemy does what is worst for you--even though it may not be rational for him to do this. This is sometimes an awful prospect and, in addition, plainly pessimistic, so one may wish to design against a "rational" rather than a malevolent enemy; but as much as possible, one should carry some insurance against the latter possibility.

Chapter 4 is mainly concerned with an over-idealized treatment of the two-sided war. There is more discussion of the conceptual side of this in Chapter 10 on Game Theory. In addition to the central ideas we have just mentioned, we discussed:

- Why it is necessary to be precise, and why this ordinarily means quantitative analysis. (pp. 1-10.)

- Examples of different kinds of models. (pp. 13-20; 46-49; and 76-79.)

- How choosing a particular model automatically narrows the range of circumstances being considered and the corresponding dangers. (pp. 21-23; 58; 59.)

- How to manipulate the model to simplify calculations. (pp. 23-25.)

- Expected value models. (pp. 26-30; 55-57.)

- Why the performance of our system was insensitive to many important parameters and that this insensitivity was not accidental, but had been achieved by proper design. (pp. 28-30.)

- The Monte Carlo technique of calculation. (pp. 48-54.)

- Examples of the ease with which models can be misused. (pp. 57-60.)

- Importance of different kinds of objectives and how they can affect the design and operation of a system. (pp. 60-72.)

- The notions of Low and High Confidence.

- The impossibility with current computing techniques of performing the necessary calculations in a routine way and the corresponding necessity for skilled approximations. The whole of this part of the book is supposed to illustrate how the analyst's intuition can be trained so that he can make reasonable approximations and how these approximations are justified. Pages 127 to 128 summarize a typical set of approximations.