U. S. AIR FORCE

PROJECT RAND

RESEARCH MEMORANDUM

AN INTRODUCTION TO SYSTEMS ANALYSIS

Malcolm W. Hoag

RM-1678

18 April 1956

Assigned to ______________________________

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

PREFACE

This material was given as an orientation talk to an Air Force group visiting RAND. Its limited purpose is an elementary exposition of some conceptual issues that are relevant for a critical understanding of Systems Analysis. Many conceptual subtleties and all issues of analytic technique are avoided.

Talks on this subject by several RAND colleagues have been drawn on freely without specific citation, especially those of C. J. Hitch (An Appreciation of Systems Analysis, The RAND Corporation, P-699), R. N. McKean, R. D. Specht, and A. J. Wohlstetter, and I am indebted especially to N. M. Kaplan for helpful criticisms.

AN INTRODUCTION TO SYSTEMS ANALYSIS

A talk on Systems Analysis ought to begin with a definition of this term. Unfortunately no precise, commonly accepted definition exists. For the moment let me say merely that by Systems Analysis we mean a systematic examination of a problem of choice in which each step of the analysis is made explicit wherever possible. Consequently we contrast Systems Analysis with a manner of reaching decisions that is largely intuitive, perhaps unsystematic, and in which much of the implicit argument remains hidden in the mind of the decision-maker or his advisor.

Systems Analysis is an outgrowth of World War II Operations Research, although it typically deals with choices that concern operations farther ahead in time, and takes a somewhat broader look at problems of military choice. Analyses at RAND, for example, typically consider what equipment to procure and develop for the Air Force of the future. Consequently the analyses that are relevant differ a good deal from the Operations Research of World War II. Dealing with the future, we are less constrained by the specific equipments and modes of operations that we happen to have at the moment. We have greater flexibility and range of choice.

This difference between a World War II kind of Operations Research and Systems Analysis is less a matter of substance than of degree. Consider, for example, the purchase of a house. If we are interested in buying an old house, an associated decision about furniture will be involved. One aspect of that decision may be the choice of a refrigerator, and we may find that a space only 30 inches wide exists in the kitchen for a refrigerator, that the house is wired only for 115 volt current, and that no gas lines are available. Consequently our choice of refrigerators is very constricted, and for that reason the problem of choice may be fairly easy. On the other hand, if we are buying a new house, one yet to be designed, our choice of a refrigerator is quite a different problem. If we sit at the drawing board with an imaginative architect, the kind of house we can have, including a refrigerator, is wide-open. We are no longer constrained to think the refrigerator must be no more than 30 inches wide, and we can consider the alternative of a gas rather than an electric refrigerator, or even an electric refrigerator that will utilize 220 volt current rather than 115 volt if that alternative is relevant. Under these circumstances, our range of choice is far broader and the number of alternatives that are relevant is consequently far greater.

Whether we speak of choosing refrigerators or military systems, some common conceptual elements are involved. Whenever we choose rationally, we try to balance the objectives we wish to attain against the costs of their attainment. In doing so certain common questions are always involved. First, what are the relevant alternatives anyway? Second, what in principle is the test of preferredness that we ought to apply in choosing among alternatives? In other words what is our criterion for choice? And, third, how do we go about the actual process of weighing objectives against costs in the selection among alternatives? To use a word that we will discuss a little later, what is our "model" of the situation? How do we go about applying it, and how do we interpret the results?

Let us turn first to the question of the alternatives that are relevant. There are two classic errors that can be made here. Each consists of going to an untenable extreme position. Consider first the one error, which is to look at an unduly restricted range of alternatives.

We should be making that mistake if, in terms of our homely illustration, we were to design a new house as if the refrigerator for it had to be constrained by all the specific limitations of an already existing old-fashioned structure. Of course, there are certain advantages in narrowing our range of choice arbitrarily in this way. It probably makes the analysis far more convenient. There are far fewer alternatives to consider, and we are less bedeviled by the problem of choosing. But we may pay a very high price for such a convenience. When we impose our constraints arbitrarily, some of the excluded alternatives may be far better than the ones left to choose from. This is a good way to get a bad refrigerator.

It is an equally good way to get a bad military system. Consider this possibility as an illustration. It is very convenient for the aircraft designer, torn by doubt about the size and other relevant characteristics of the payload that his airplane will carry many years hence, simply to assume that the physical configuration of a bomb will then be pretty much what it is today. Similarly it may be very convenient for the designer of bombs to assume that the bomb bays of the future will be of the same capacity as those of current models. But look what can result from such shared convenience. The airplane designer may believe that his airplane must be big in order to be able to carry a bomb as big as the current ones. In turn the landing space, whether it be fixed or floating, may have to be big. We end up with a combination of big bomb, big airplane, and big landing space. Yet it is quite possible that a more efficient system on the whole could be devised by combining a bomb small in size if not in bang with a small airplane operating from existing small airfields or carriers. Maybe such a better system would not be possible, but the point we want to make here is simply that whether it is possible or not can only be ascertained by looking at it. For that matter, an analysis that proceeded in such a narrow way would not look either at alternative systems that employed bombs much bigger in size than current models. If we look at an unduly narrow range of alternatives to begin with, we may never even pose the question in the right way. And getting a neat answer to the wrong question may be worse than an incomplete answer to the right question.

We can avoid the error of looking at an unduly narrow range of alternatives by specifying that we shall look only at broad systems for a comparison, and that all of the alternatives that are relevant in a broad system will be considered. In terms of the above illustration, we can try to look at all possible combinations of bombs of the future with airplanes and bases of the future. The great difficulty which this procedure raises, of course, is that of workability. The more inclusive the system, the more alternatives that are relevant, and the greater the difficulties of comparison.

Too few people appreciate the manifold alternatives that are available when one looks at a relatively unconstrained situation. Consider the simplest sort of arithmetic illustration. In asking what bombing airplane we ought to design for the future, let us concentrate arbitrarily for the moment on just three obviously pertinent characteristics -- range, speed and payload. More of any one of these is obviously a good thing in itself. We should like to fly farther, faster, and with a bigger payload. But, of course, we achieve none of these without cost. Indeed, in relevant cases we probably can achieve more of any one of these characteristics of performance only at the cost of the other two, let alone at the cost of still other considerations. If we want to fly faster, we must accept some range penalty, and so on. But suppose we have just these three characteristics to consider, and suppose, again quite arbitrarily, that to each of them we assign only a "high" or a "low" value. We can have a fast airplane or one not so fast, and so on. Now how many alternative combinations of these three characteristics do we have? We can answer this question by using the diagram given below in Chart I. If we take any one of these characteristics as a beginning point, it can assume either of its two values -- high or low. Each of these values in turn can be matched up with either the high or the low value of the second characteristic, so that we have four distinct combinations. Each of these combinations in turn can then be matched with either the high or the low value of the third characteristic. What we end up with are (2)(2)(2) or 8 combinations.

Chart I

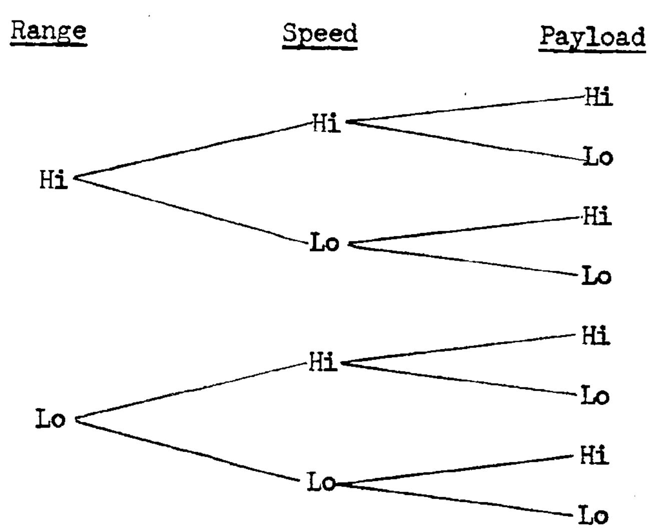

The number of combinations of interest is not two multiplied by three, i.e., six, but rather two to the third power, which equals eight.

The number of combinations increases exponentially, not by simple multiplication. Now that does not matter much when the difference is only between six and eight. But clearly we are interested in many more characteristics of our airplane. What about crew space, landing facilities, and electronic gear for instance? Yet if we include only ten characteristics of interest, and still restrict ourselves to but two values of each, the number of relevant combinations is $(2)^{10}$ or about 1,000. The difference between 6 and 8 combinations is small, but the difference between 20 (2 times 10) and 1,000 is large. Typically when we start with clean drawing boards and are free to consider all alternatives of interest in a problem of the future, there are many millions of possible alternatives. The problem is one of choosing sensibly from among that large number. Given that kind of a situation, it is manifestly absurd for anyone to say that we should develop all alternatives in a hardware form, see how they perform, and then choose among them. It is imperative to find a cheaper way to compare them. The bulk of the alternatives must be excluded by a comparison that utilizes few, if any, tools beyond pencils, paper, and discriminating thought. Given that this is the situation, there is surely a case for making the choice on paper as carefully and as ably as we can. Hence the desirability of Systems Analysis.

You notice that all I have done in discussing the range of alternatives is to point up the error of going to either extreme. The unduly narrow analysis may exclude the really interesting alternatives. The impossibly broad comparison is simply out of the question. Nobody but a fool or a charlatan ever pretends to look at all the universe and to resolve all problems at the same time, even though in principle those problems are interdependent. In a general discussion there is very little that can be said positively about how big the context of a problem ought to be. It depends on the particular problem at issue. All one can say is that the analyst must be very careful to make his analysis neither too small nor too large, which is only to say that judgment is essential at this very early stage.

The problem of the criterion, or test of preferredness, that ought to be used in a particular problem is always with us. To take one aspect of our military problem as an example, let us consider the question of the preferred engine for the bomber of the future. Suppose, again arbitrarily, that only turbo-jet engines are at issue. One possible criterion for choosing among alternative engines is that of pounds of thrust per pound of engine weight. Clearly greater thrust is desirable; equally clearly, light engines are desirable. We can balance the one characteristic against the other for the alternative engines that we think we could have. The result of a comparison might appear as in Chart II below.

Chart II

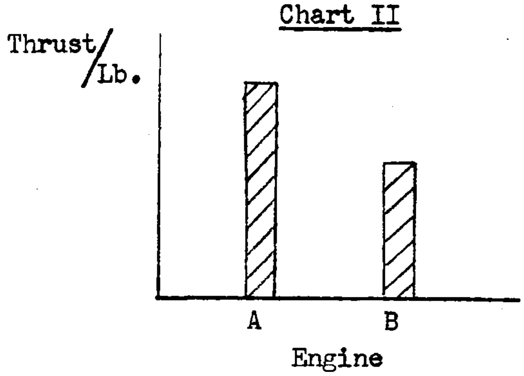

On the basis of this comparison alone we should prefer Engine A to Engine B. But any sensible engineer will immediately object that this is too simple a criterion for engine choice. It leaves out of account many important characteristics, and there is no assurance that these other characteristics will influence the decision in the same direction as the two characteristics of thrust and engine weight that are considered. An engineer will immediately ask, "Is engine A more durable than B? Is its specific fuel consumption lower? Is it cheaper to produce?" Certainly these are relevant questions, and they do not answer themselves. It is possible that Engine B is decidedly superior to A with respect to one or more of these additional characteristics, and there are undoubtedly other relevant characteristics as well.

Moreover, we may note in passing that there is one other question that certainly ought to be asked. What size of engine in absolute terms results from following such a criterion to the bitter end? It is possible that we would get a 60,000 pound thrust engine, one so big that we might not be interested in it at all. The use of a ratio as a criterion for choice yields no assurance that the absolute scale of the operation in question will be an interesting or relevant one.

Given the difficulties that we can get into with an obviously incomplete criterion, what are we going to do about it? The only way to resolve the question is to look at the comparison in terms of the context of a broader system where we can bring a more complete criterion to bear upon the problem. For example, we can consider the context of a strategic bombing system as a whole. Suppose we take a "reasonable" list of Soviet targets for a strategic bombing campaign, and stipulate that the objective of all of the systems to be compared is to acquire as great a capability of destroying those targets as possible for a given total budget. (You will notice that I avoid all question of what constitutes a "reasonable" target.) What system that can be devised will have the greatest capability given the budget constraint?

As we experiment in our paper comparison with different bombing systems, we are free to choose among engine types. We shall try to choose that engine which, when married to all the other desirable components of the bombing system, is consistent with the greatest combat capability. The issue between engine A and engine B will now tend to be resolved. If we only listed the technical characteristics of engines A and B, it is likely that we would have two columns: one column of characteristics in which engine A was the superior of engine B, and a second column of different characteristics in which engine B was superior to engine A. So expressed, there is no common denominator in terms of which of these different characteristics can be compared. Such a listing would suffice for making a choice only if one engine uniformly dominated the other, that is, was better in each characteristic. One just does not expect to find that sort of dominance in an interesting problem. Interesting problems are typically not that easy. Therefore, in order to choose between the two engines we must find a way to reduce their diverse merits and demerits to a common basis. The context of the strategic bombing system as a whole makes it possible for us to find that common denominator. Each merit of a particular engine type tends to have its impact in terms of a contribution toward greater target capability. Conversely each demerit tends to have its effect in terms of decreased capability. The net effect of the merits and demerits of each will result in a final contribution to total system capability. Their differences can be expressed in terms of the common denominator of targets, and which engine is better can be established given a budget constraint.

How can we determine the appropriate scale of the bombing campaign at issue? How big a budget should we impose as a constraint? It may well be that we cannot answer this question with any precision. In that case all that we can do is to try several budget levels and see what happens at each level. If we are lucky we find that our preference for one engine type over the other does not depend upon the budget level. That is, choice between the two is not sensitive to variation in the scale of the campaign. If so, we can make a strong recommendation for that engine.

The criterion in our bombing campaign comparison is expressed in terms of maximum target capability for a given total budget. The criterion could equally well be expressed the other way around in terms of minimum total budget for a given target capability. At any particular scale these two expressions normally amount to the same thing.* If we express the criterion the one way, we get a result like that shown in

---

* The abnormal case where these expressions would not amount to the same thing requires that one alternative be more expensive than another at low objective levels, less expensive at high objective levels, and have the same or lower total costs associated with higher objective levels, not just lower average costs per unit of the objective. This case is unlikely in practice. However, for the curious:

Suppose Alternative A can get 5 targets at a minimum cost of $50 m., but Alternative B can get 5 targets only at a minimum cost of $55 m. If we are only interested in 5 targets, A is unmistakeably our choice. But suppose we are interested in more targets, and B could get us 20 targets at a total cost of $50 m. (a lower total cost than that associated with 5 targets for B!) while A could get us 20 targets only at a cost of $60 m. In this case, if we express our criterion in terms of maximum capability for a $50 m. budget, we get an answer of B and 20 targets. All is well. Should we express our criterion in terms of minimum cost for 5 targets, however, we get an answer of A and the same cost of $50 million but a lower capability, clearly an inferior solution. Here it makes a difference which way we express the criterion.

The trouble arises in this unusual case because the meaning of "at any particular scale" is ambiguous. Because total costs are not always rising with higher objective levels, more than one objective level is associated with any given cost level. No ambiguity arises in the usual case illustrated in Part (2) of Chart IV, and we can identify a particular scale by reference either to a cost or an objective level.

Part (1) of Chart III; if the other way, we get a result like Part (2) of Chart III. In either case the margin of superiority of systems employing engine B over systems employing engine A is much the same. The result of our comparison in the context of a broader system is to correct the misleading impression of a technical superiority of engine A over B that was generated by the overly simple criterion of maximum thrust per pound of engine weight (Chart II).

The comparison in Chart III is in terms of a scale of a $1 billion budget (Part 1) or 100 targets (Part 2). As we try different scales of effort, notice what happens. At each scale that we try, say $1 billion, $2 billion, and so on, we shall try to establish that system with maximum target capability. Once that system has been found, its cost can be transposed to a different chart that will graph the relationship between the number of targets that can be destroyed and the total cost necessary for their destruction. We end up with the chart in Part (2) of Chart IV.

Incidentally, such a chart serves to demonstrate the fallacy of one commonly encountered criterion. What does it mean if somebody says the criterion for this job is to get the most for the least? The "least" in this case means lower dollar cost, which means we move to the left in Part (2) of Chart IV. But getting the "most" means that we destroy the most targets, which means that we go up. But we cannot go up and to the left simultaneously in this chart. Higher capability means higher cost once we have determined the efficient systems to be employed at various scales of effort. Such a criterion involves us in a simple verbal contradiction. To convert it into a meaningful proposition, we must rephrase it by saying, "Get the most for a stipulated cost; or, alternatively, for a stipulated objective get the lowest cost that you can."

Chart III

Alternative Criterion Wordings

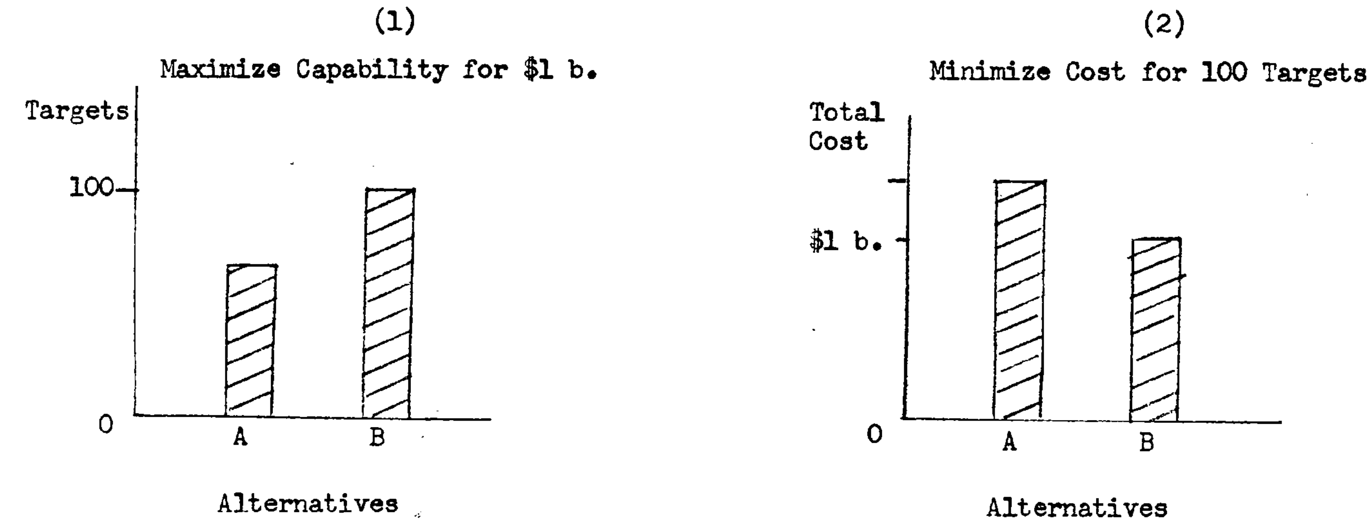

Chart IV

From One Scale to Many

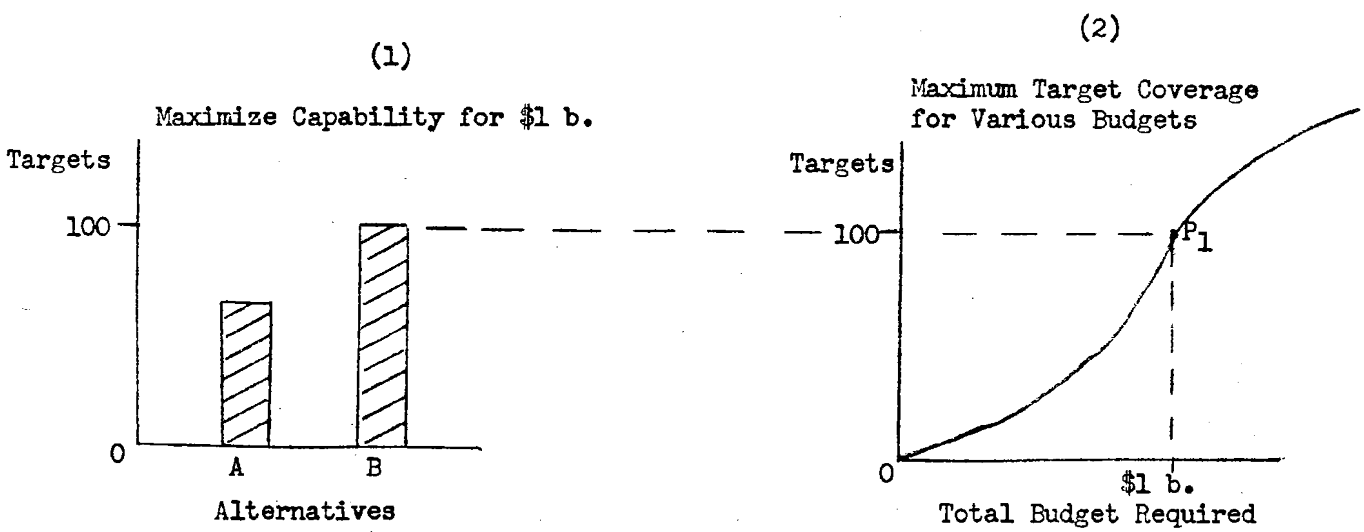

So far I have ignored the vital matters of how you go about the actual process of comparison in a Systems Analysis. We must simulate a real life performance by alternative systems, and naturally we want to make the simulation as realistic as possible for the important factors at issue. I stress the word "important." It is very clear in any analysis that we cannot take everything into account. Something must be left out in order to make the analysis workable. Given that we have to eliminate something, we shall try to eliminate the less important.

Sometimes the problems of our military analyses are compared to the problem of the owner of a racing stable who wants to win a horse race to be run many years hence, on a track not yet built, between horses not yet born. Everybody would agree that a prediction of the kind of grass that would grow in the infield of the race track is quite unimportant in the analysis, and need not be included. But somebody would certainly quarrel with a failure to consider the composition of the running surface itself. Race track fanciers tell me that some race tracks are faster than others, and that some horses perform relatively much better on slow tracks than they do on fast tracks. Again we only point to something at once important and difficult; we do not tell you how to do it. How do you decide which factors to suppress and which ones to include? And yet an analysis not only of necessity suppresses some things; it aggregates many of the obviously important things that remain. A multiplicity of small things must be aggregated into a small number of manageable big things. Only when we have done this will the problem of comparison be reduced to manageable dimensions. Consequently we shall simulate real life incompletely by concentrating only on the things that we think are important and abstracting from the unimportant. In other words, we shall create a "model."

Most of the actual work in any analysis will deal, of course, with the choice of factors to include in the model and their quantitative estimation. The great importance of careful, skilled estimation is obvious. If we want to build the best home, we ask how alternative building materials compare in strength, weight, aesthetic appeal, and cost. If we want to build the best bomber, we ask what performance characteristics are attainable within the expected state of the art in aircraft construction, what relevant trade-offs can be made among those characteristics, and what defenses future bombers may have to face. There are many more relevant questions that could be listed in either case. All must be answered as best they can be in order to make a comparison possible, and the better the answers, the better the analysis as a whole. There is little more to say at a level of general principle except that it is important that the various degrees of confidence with which particular factors are estimated should be expressed, and the variance attached to particular estimates is important as well as the best guesses about the expected magnitudes.

Given a common denominator with which to compare physically diverse things, all or a good part of the model will typically be quantitative in character. Hence the talents of the mathematician, of value elsewhere in the analysis, are indispensable at this stage. We turn to a mathematician not so much for his skill in performing mass computations as for his ingenuity in escaping them. With judgment, one may be able to boil down an unworkably large number of alternatives to a number which, while formidably large, is manageable. With ingenuity, one may be able to substitute simplified for intricate mathematical manipulation. When all of the honest tricks of the trade that permit such reduction have been utilized, we have the electronic calculator to fall back upon to handle masses of computations. Improvements in the art of mathematical manipulation and the development of the electronic calculator have contributed greatly to the potential for useful systems analyses.

So far we have talked about several common conceptual elements that enter any problem of choice. But behind all of them is one troublesome element of which you are all aware. When we talk of the future we are speaking of a future largely unknown. Any analysis of current decisions that will bind us in the future is plagued by the difficulties raised by uncertainty. For military studies, what will the enemy be like in the future? How good are our forecasts of the technological improvements that will be possible if we will only give the go-ahead signal now to certain development and production programs? And in the event our military devices are actually employed in the future, what chance elements will help to determine their success? How lucky or unlucky will we be in the particular situations?

We are uncertain about all of these things, and our uncertainty gives us a two-fold problem. Because we are so uncertain about the magnitudes of many of the important factors, we don't know how to push a comparison between alternatives to a definitive conclusion even when we have a definite criterion for choice between them. As if that were not bad enough, the formulation of a criterion for choice in uncertain situations may itself be elusive. To take the simplest sort of example, different people may give different answers to the question of whether they would prefer a certainty of $5 to one chance in twenty of getting $100. Yet the mathematical expectation of gain is the same in both cases. We cannot fix one definite criterion for choice between these two that will apply to all individuals.

Granted that uncertainty tends to be pervasive, what can we do about it in Systems Analysis? Well, first we can do the same sort of thing we did when faced with uncertainty about the budget level in our example. Once we have a particular result based upon one set of assumptions, that result can be tested for its sensitivity to changes in many of the factors that enter into the analysis. We can try an altered scale of effort by assuming a different budget, we can double the strength of the estimated enemy defenses, and so on, and see what such tests do to our choices. One essential of most systems analyses is a battery of sensitivity tests that show how the results depend upon different factors. If we are lucky, we emerge with a result that is comparatively insensitive to a wide range of change in many critical elements of the analysis. If we are unlucky, we emerge with at best equivocal recommendations, although the factors upon which choice critically depends, and the particular uncertainties that it is important to try to resolve, will have been clarified.

For dealing with uncertainty, some kinds of statistical uncertainty can be treated explicitly in the model. This possibility raises technical matters that I do not propose to go into here beyond noting that the criterion can be expressed in a somewhat different form. Instead of seeking the minimum cost system among systems all of which have an expectation of attaining a certain capability, the criterion can be rephrased, for example, in terms of a confidence level: "Choose the least costly of the systems that have a 90 per cent assurance of being able to do a stipulated job." Given this different sort of criterion, different model manipulations may be called for, and one may get a somewhat different solution.

Finally, and probably most important of all, we can try to build into our analyses a provision for hedging our bets. We can try to protect some of the bets we make upon an uncertain future by making others in opposed directions. In that way we may not be caught out too badly by particular shifts of fortune. For example, if we are making recommendations for novel types of equipment at quite an early stage of development, we may want to recommend multiple programs: not "A is preferred to B," but "develop at least both A and B from among a much larger number of possible alternatives." If both are developed, we are much less exposed to catastrophe than if we had plumped all out for only A or B. The chance that both will be failures is perhaps much lower than that either will be a failure if we concentrate exclusively upon it. To revert to our future horse-race analogy, one may want to analyze carefully in choosing preferred breeding stock, but prefer to seek more than one colt from more than one set of parents. But not, of course, without limit. Stud fees and the prices of brood mares are high for blood lines of proved merit. And modern hardware development for military application is not cheap either.

The principle of hedging can be extended all the way through to force composition and employment decisions. One may argue cogently for systems of mixed weapon types on the basis that we are less uncertain about getting tolerable results than in the case of a pure system of one type, even in cases where our best guess of probable results tends to push us toward a pure system. But the cost of this type of hedging rises as we move from the early stages of research and development toward procurement and use, and we can afford less of it. Moreover, as we move from early to late stages of development the burden of uncertainty lessens, and hedging loses some of its great value. Incidentally, this consideration suggests that we supplement our hedges by not making particular decisions prematurely. If there is no necessity to choose an engine now because the weapon program at issue is bound to be held up for a long period of time anyway by the need to develop a guidance system, it may be very desirable to defer our engine choice to a time when our decision can be a better one.

You notice that none of these comments tells you precisely what to do in particular situations. It is very clear that Systems Analysis as currently practiced, and probably as practiced in the future, is much more an art than a science. To be sure the analyst ought to employ whatever available scientific tools are appropriate at particular places in the analysis, but his operations as a whole are not characteristic of "Science" with a capital "S." We have stressed the part that judgment plays at each stage of a Systems Analysis. What scope should the analysis take anyway; what is the appropriate balance that escapes being overambitiously big or unduly constricted? What criterion ought to be applied? What factors can be suppressed in the analysis? Of those that remain, what degree of aggregation of particular factors is permissible? How are they to be manipulated? How is the result to be interpreted?

Moreover, our very division of Systems Analysis into "stages" has too much of a cut-and-dried flavor, although that division is useful in order to talk about some conceptual issues one at a time. In practice nobody is likely to give you a problem neatly and correctly formulated, with the identification of relevant alternatives a matter for simple enumeration and the other stages to follow. One of our senior practitioners of the art of Systems Analysis insists that the term itself is unfortunate and should be replaced by Systems Design. In his view the perception of the problem and the finding of relevant alternatives are the key creative parts of a study, and the use of the term "Design" is meant to highlight those parts. The imagination to perceive an important problem before others have become aware of its imminence, and the inventiveness to devise novel alternatives for its solution, are perhaps the most important qualities the systems analyst needs. And, like the sound judgment also required, these qualities cannot be taught in terms of general principles or set rules. One can only repeat admonitions: "Systems Analysis is no substitute for sense," and "There are neither prescriptions nor substitutes for ingenuity in analysis."

Sometimes when stress is properly put upon all the difficulties that preclude a definitive analysis of decision problems in terms of a full-fledged scientific discipline, one encounters the following sort of reaction to a study: "There are so many imponderables in the situation, so many uncertainties, that I just have no faith in the results that emerge from all this mass of data and discussion. I prefer to avoid it, and therefore I avoid Systems Analysis." There is a direct answer to this type of understandable objection. The difficulties are intrinsic in the nature of the problems that are studied. They are not difficulties that are grafted onto the problems just because particular techniques are brought to bear on them. The same burden of uncertainty must be borne by anyone who tries rationally to arrive at decisions. The man who solves complex problems in the span of five minutes on the intuitive basis of "sound" military or business judgment and experience, and that is certainly possible, may then forget about it and sleep easily at night. In contrast, the man who has worked two years on an involved Systems Analysis of the same problem may sleep badly, but only because he has become acutely aware of all the pitfalls in the problem.

At RAND we tend to have faith that Systems Analysis can do a much better job on many problems than conventional staff work in the military organization, or, for that matter, elsewhere in government or industry. If you take bright people who are scientifically trained, who are detached from the particular problem and consequently have no bias, and give them ample time and opportunity to bring their combined diverse experience to bear upon the problems, they certainly ought to be able to do considerably better than one harassed Indian who is given, say, three months to solve an impossibly big problem with little assistance. And they ought to do better even when the comparison is a fairer one.

One should not overstress the difficulties of finding the best solution to complex problems by analysis, even though the difficulties are great. To be sure, many problems by their very nature are so complex that one would never be confident that the best result had been found. At the same time very useful results may be achieved. In particular, one may tackle a problem in order to find a better rather than the best solution. Instead of starting with a clean drawing board, looking at all the relevant alternatives that can be discovered, and endeavoring to find the best, one may take as a starting point the existing program or pattern of decisions that, unless shortly altered, will bind the future. Taking that as a starting point, in what respects does the program seem to be demonstrably bad? If such respects can be found, one can proceed to the demonstration that they are bad. One seeks an alternative program that can be expected to do better in the likely contingency, and yet be better adapted to meet unlikely but possible contingencies. Probably we would be better advised in many problems to tackle the design of better systems rather than the abstract analysis of the best. If we did, we might be able to make useful recommendations that would otherwise escape us, for we should bog down in inextricable difficulties if we pursued the best of all possible worlds. At the worst, definitive recommendations may not be possible, but analysis will at least make more explicit for the decision-maker many elements pertinent to his decision of which he is imperfectly aware. At the best, one may do far better. Demonstrably important and secure recommendations may be produced.

It is only fair to add that while Systems Analysis can and ought to be much better than conventional staff work, it can be worse. Many systems analyses share an impressive facade. The technical discussion is long and complicated, the charts are elegant, the mathematical appendices are formidable, and there is great display of technical jargon and virtuosity. Such a facade can reveal very good analyses or it can conceal very bad ones. If it conceals bad ones, it may be worse than no analysis at all. The difficulties of understanding the technical display may impede rather than facilitate relevant criticism. The whole elaborate structure can conceal a big butch or a big bias in terms of arbitrarily excluded alternatives, an improper criterion, bad execution of the analysis, or poor interpretation of the results. But all tools can be used or abused. The possibility of misuse need not and should not deny us the opportunity to make good use of the available tools. After all, "The substitute for bad analysis is good analysis." How else do you know it is bad?

MWH:dmw