14-p

AD613003

UNCERTAINTIES IN OPERATIONS RESEARCH

Charles Hitch
Economics Division
The RAND Corporation

P-1959

25 April 1960

DISCLAIMER NOTICE

THIS DOCUMENT IS BEST QUALITY AVAILABLE. THE COPY FURNISHED TO DTIC CONTAINED A SIGNIFICANT NUMBER OF PAGES WHICH DO NOT REPRODUCE LEGIBLY.

UNCERTAINTIES IN OPERATIONS RESEARCH

Charles Hitch
Economics Division
The RAND Corporation

P-1959

25 April 1960

UNCERTAINTIES IN OPERATIONS RESEARCH

Charles Hitch

What I want to do this evening is to talk in a general, but I hope
constructive, way about some of operations research's most intractable
problems -- those associated with uncertainties, and especially with those
uncertainties tinged with "game" elements. My own operations research
experience is pretty much restricted to the domain of military problems,
but for an examination of the implications of uncertainty, this doesn't
seem to me to be too important. For the kinds of uncertainties that we
encounter in military problems have their counterparts and analogues in
business and in everyday life.

In fact, no characteristic of decision making is as pervasive as un-
certainty. When, as operations researchers, to simplify a first cut at an
analysis, we assume that the situation can be described by certainty equiva-
lents, we may be doing violence to the facts and indeed the violence may be
so grievous as to falsify the problem and give us a nonsense solution. How,
for example, can we help the military make development decisions -- decisions
about which aircraft or missiles to develop -- when the essence of the prob-
lem is that no one can predict with accuracy how long it will take to develop
any of the competing equipments, or to get them operational, how much they
will cost, what their performance will be, or what the world will be like
at whatever uncertain future date turns out to be relevant (if indeed, the
world still exists then)? When I say "cannot predict with accuracy" I am
not exaggerating. We find that typically, for example, the production costs
of new equipment are underestimated in the early stages of development by

factors of two to twenty (not 2 to 20 _per cent_, but _factors_ of two to twenty.) Why they are always underestimated, never overestimated, I leave to your fertile imaginations.

A business is confronted with uncertainties of similar type and order when it contemplates diversification. So is a household when it tries to decide whether to throw away empty boxes that may conceivably be useful at some future time, or let them continue to clutter up the garage. So is the same household when it tries to decide what in the way of clothes to bring to New York for an occasion like this. Who knows what the evening temperature will turn out to be? Who knows what other people will be wearing?

The really remarkable thing is that while we operations researchers have no good general rules for how people ought to decide when confronted with such uncertainties, men nevertheless do make decisions. Here I use the word "men" advisedly, for women appear to have much greater difficulty. There is a real difference between the sexes. Women apparently appreciate the absence of good theory, of accepted principles of rationality, and are in consequence paralyzed into indecisiveness like a too sophisticated and too conscientious operations researcher. Men, on the other hand, don't appear to be bothered. They decide impulsively, like the uncritical operations researcher who mistakes guesses about the future for facts, or the operations research charlatan who conceals the uncertainties from his customer for fear his recommendations won't be accepted.

High speed electronic computers and data processing equipment aren't of too much assistance in coming to grips with the problem of how one ought to decide. My distinguished predecessor, Bernard Koopman, has already dealt adequately with what he calls the fallacy of "mechanitis." Of course,

computers do help in a mechanical way. If there are a dozen interesting future contingencies, high speed computers will facilitate the computation of outcomes in all twelve. They facilitate Monte Carlo calculations which display a range of possible outcomes. But they still leave you with the fundamental problem: some courses of action will seem to be better in some contingencies, others in other contingencies. Which is optimal when you can't know which contingency will occur?

Worse, in almost all our interesting problems, there are processes whose results we cannot even predict as probability distributions. There is no sharp line of demarcation between insurable risks, like rain and natural death, where there is an analyzed mass of data on similar occurrences to provide an objective basis for calculating a probability distribution; and genuine uncertainties, where such data are simply lacking. Lloyds of London doesn't draw the line at the same point as the typical American insurance company, and the poor operations researcher must frequently make a stab at situations that Lloyds of London would shun. How long will it take us to develop equipment capable of getting a man to the moon and back -- say, as a function of the rate of effort on the development? Most of us have some feel for how we would go about making such an estimate, but no two of us would get the same answer except by grotesque coincidence, and only the most foolhardy and foolish would have any confidence in his answer. Of course, bright people have been confident about similar predictions in the past (or have pretended to be for some ulterior or praiseworthy purpose). But the record shows no correlation between confidence and accuracy. And if predictions are uncertain in the domain of mechanics, contemplate some less well understood fields. For example, let's think of estimating the time required

to develop a cure for cancer as a function of the rate of effort. I'd decide
pretty quickly that I could put my own time to better use. The world is
going to have to get along somehow without the answer to that one. It was
E. Bright Wilson who said: "Many scientists owe their greatness not to their
skill in solving problems but to their wisdom in choosing them." (It is
perhaps significant that the great scientists to whom he was referring have
not rushed into operations research.)

Even with hindsight it's sometimes pretty hard to trace causal relations
with a time dimension. Two veterans of World War I were recently talking
about their experiences in France in 1918. "Remember," one of them said to
the other, "that camp back of Verdun where they fed us all that saltpetre?
Well, you know, I think it's beginning to take effect."

And most perplexing of all, beyond insurable risk and genuine uncer-
tainty, there are the problems of intelligent opposition. Of malevolent
opposition. And the similar and equally perplexing problems, which we also
classify as "game" elements, of cooperation, of alliances, of substantial,
but nevertheless incomplete community of interest. How can we predict with
any confidence what an intelligent (or unintelligent) enemy or ally will do?
And how can we make decisions when a good decision depends on what enemies
and allies will do?

One thing is certain -- the operations researcher ignores all these
uncertainties at his peril. Suppose that we are uncertain about 10 factors
in an analysis (factors like -- will we be able to use overseas bases, will
the enemy have interceptors operating over X,000 feet), and suppose that we
think our best guess in each case has a 60 per cent chance of being right
(we'd be lucky indeed to have such good information.) If we confined our

analysis to the ten best guesses we would be ignoring a set of future out-
comes which, taken together, have a 99-1/2 per cent probability of occurring.

After wrestling with problems of uncertainty for a number of years I
have come, hand in hand with a number of my colleagues, to some unoriginal
but I believe true and important conclusions. We are skeptical of general
purpose solutions to the problems of making good decisions for uncertain
contingencies, and are convinced that a shift in emphasis is called for in
our approach to such problems. A shift from searching for the best way to
choose between two contingently unsatisfactory answers to searching for a
better answer. From a search for a better decision rule to a search for a
better system. From sophistication in judgment to ingenuity in design.

Let me touch lightly on the unsatisfactory character of the general
approaches to decision making under conditions of uncertainty. The reasons
for lack of satisfaction are fairly familiar.

First, the maximization of expected or "average" outcomes, values, or
"utilities." The troubles here are formidable. The calculation of an
expected value requires knowledge of two things: of the probability distri-
bution of expected events, and of the values of the outcomes associated with
each. But the probability distribution is known only for insurable risks;
i.e., not often in our business. Where there is true uncertainty we must
either try by research to reduce the uncertainty to something approaching
insurable risk status, or fall back on the thin reed of subjective prob-
ability.

Similar or worse difficulties confront us in establishing the values
or utilities of the outcomes. In playing against nature the maximization of
expected utility has a very strong intuitive appeal. But it is a rare problem

indeed in which we have any generally acceptable means of measuring utilities. So we tend to fall back on the expected outcomes themselves -- in physical rather than value terms. Instead of maximizing the expected utility or military worth of the targets our system will destroy, we maximize some physical quantity like the expected number of targets destroyed. It is so much easier that way. But no one who has thought for as much as five minutes about this procedure could defend it. Typically it leads to a reckless choice of strategies, for it ignores the importance, which is frequently crucial, of diminishing marginal utility. In general the more we have of some physical good -- like dead Russian targets, or speeches at a dinner -- the less intensely we desire one additional unit. Cases even exist where the extra utility associated with an additional unit becomes negative. But when we maximize expected outcomes we simply assume away this fact of life. We value equally a system which we expect to kill the enemy twice over half the time, and to be completely impotent the other half, and a system which kills him dead for certain, but only once.

The other trouble with simply maximizing expected outcomes or utilities or anything else is that it leaves the game elements out of the problem -- at least explicitly. This may be all right if you are playing against nature, but who plays against nature? You may think you are when you play roulette in Las Vegas, and you are probably right in thinking so, but I assure you that some who have tried are convinced that a malevolent intelligence has been working against them. Nature may be neutral, but sometimes doesn't seem so. Ponder for a moment the experience of the Barbadoes bricklayer who wrote the following letter requesting sick leave of his employer:

"Respected sir, when I got to the building, I found that the

hurricane had knocked some bricks off the top. So I rigged up a beam with a pulley at the top of the building and hoisted up a couple of barrels full of bricks. When I had fixed the building, there was a lot of bricks left over.

"I hoisted the barrel back up again and secured the line at the bottom, and then went up and filled the barrel with extra bricks. Then I went to the bottom and cast off the line.

"Unfortunately, the barrel of bricks was heavier than I was and before I knew what was happening the barrel started down, jerking me off the ground. I decided to hang on and halfway up I met the barrel coming down and received a severe blow on the shoulder.

"I then continued to the top, banging my head against the beam and getting my finger jammed in the pulley. When the barrel hit the ground it bursted its bottom, allowing all the bricks to spill out.

"I was heavier than the barrel and so started down again at high speed. Halfway down, I met the barrel coming up and received severe injuries to my shins. When I hit the ground I landed on the bricks, getting several painful cuts from the sharp edges.

"At this point I must have lost my presence of mind, because I let go to the line. The barrel then came down giving me another heavy blow on the head and putting me in the hospital.

"I respectfully request sick leave."

Whether or not the Barbadoes bricklayer was playing against a malevolent enemy, military operations researchers almost always are. And so are

industrial operations researchers. And where either enemies or allies are involved, the whole notion of independently given probabilities makes much less sense, and so therefore does the calculation of expected utilities or outcomes.

Game theory, of course, was invented to fill precisely this breach, and there is no doubt that it has helped to do so. There is also no doubt that the breach is still far from satisfactorily filled -- that game theory must be developed considerably in several different directions before it becomes a very practical or useful tool for aiding decision makers.

The difficulty is a fundamental one. The only kind of game situation for which we have operational solutions is the two-person zero-sum variety. And very rarely in the real world does the military or industrial operations researcher encounter any situation approximating two-person zero-sum. It is apparent that there is nothing zero-sum about war, expecially the modern thermonuclear kind. What Russia loses we do not necessarily gain; in fact, we and the Russians have a strong common interest in eschewing certain strategies, like end-of-the-world machines. Similarly in business not only are many more than two "persons" typically involved, but our relations with none of them are those of pure conflict (i.e., "zero-sum"). A company will compete with other companies in the same industry, but not in a your-loss-is-my-gain manner; there is broad mutuality of interest too. And its relations with customers, with suppliers, and with labor are likewise extremely complex mixtures of conflict and cooperation.

It is frequently argued that max-min, the solution for the two-person zero-sum game, is unduly conservative. Whether it is conservative when the game is really two-person zero-sum I'm not sure: perhaps not. But in almost

all real games or circumstances, which are not two-person zero-sum, a max-
min solution is either conservative because it greatly limits your opportun-
ities to capitalize on the enemy's mistakes -- or worse, it misses some
vital element of the problem, like an important mutuality of interest.
Clearly in playing against nature it can be outrageously conservative to
concentrate exclusively on minimizing the worst that malevolent nature can
do to you. Nature isn't that malevolent.

In playing against the Russians or any enemy in this thermonuclear age,
max-min is worse than conservative; it completely ignores the strategy which
almost all of us intuitively prefer, the strategy of deterrence, the strategy
of preventing the war from happening at all. We attach very high positive
value to this outcome of no thermonuclear war. Do the Russians therefore
attach a high negative value to it, as the zero-sum assumption postulates?
Not if they are as smart in choosing strategies as they seem to be in develop-
ing rockets. Instead of max-min-ing, instead of concentrating exclusively
on minimizing the worst the Russians can do to us, it may well pay to devote
some of our energies to promoting an outcome which both we and the Russians
would value highly. Mutual deterrence is preferable to mutual near-annihila-
tion, even if we have made ours not quite so near by effective max-min-ing.

Unfortunately, such partial mutuality of interest can seldom be accommo-
dated by marginal adjustments in a max-min solution. It is likely to alter
the whole character of the problem, and the solution. For example, one of
the universal characteristics of good two-person zero-sum play is the com-
plete absence of communication between the players. Each chooses his stra-
tegy in ignorance of his opponent's choice. In the case of a pure max-min
strategy it does the opponent no good to learn your strategy. In cases

where he can profit by learning, y   make it doubly impossible for him to learn by not letting yourself know. You mix your strategies, and pick one for each play at the last moment by spinning a wheel of chance. You establish your command post in Las Vegas.

Contrast the real-life problems of playing with the Russians the kind of non-zero-sum game in which there is mixed conflict and mutuality of interest. Communication, explicit or implicit, becomes the essence of the problem. If we are to prevent the Russians from max-min-ing (which might well involve their attacking us), we have to reveal to them that our payoff matrix is not the obverse of theirs, and what it is. If we have a clever and effective technique for protecting our retaliatory force from surprise attack, we must somehow, without giving them information that would seriously reduce its effectiveness, make them know that it is effective. You don't keep your ability to retaliate a secret. If we sense that the Russians mistakenly fear a surprise attack by us, and therefore may try to preempt us, our most important problem is somehow to convey to them <u>and make them believe</u> that we have no such intention. These problems of communicating are perhaps most difficult and most critical in finding strategies to fight limited wars that will, among other desirable things, keep them limited.

I am not just saying that someone must find ways of solving n-person variable-sum games. I doubt that pat solutions will be found -- at least in our generation. What is needed, I think, is more <u>attention</u> to certain aspects of n-person and non-zero-sum games. And not only by professional mathematicians. As Tom Schelling has recently been stressing, and as von Neumann well understood, the outcome in n-person games is likely to depend

critically upon psychological and sociological factors which we cannot yet describe, explain, or manipulate in rigorous, mathematical language.

Let me pass quickly over another technique for dealing with game uncertainties; namely, gaming. Gaming is not tied to two-person zero-sum assumptions, and is therefore a broader and more flexible device than two-person zero-sum game theory. I am all for gaming and other forms of simulation. I think they are invaluable for teaching purposes -- and of course they have long been so used by the military. I think that free play games can also be very useful in providing "insights" to players and observers, in suggesting important factors that might otherwise be overlooked, etc. But as a problem-solving device gaming in any of its numerous forms has severe limitations. Many of these stem from the fact that gaming isn't analytic. You can't do much testing for sensitivity. You can't afford to. You never know why the games came out as they did. It may have been because you played too few times and got a badly biased sample of results. Or because some of the players were too stupid for their real-life counterparts -- or too clever. Or because the players were using wrong pay-off functions. Or because they were communicating too much, or too little. I have yet to see a case where a hard problem was convincingly solved by a game. I have, on the other hand, seen a good many cases of gaming that I thought were worth the cost -- which is always high.

So what does the poor operations researcher do? Here he is, faced by his fundamental difficulty. The future is uncertain. Nature is unpredictable, and enemies and allies are even more so. He has no good general purpose technique, neither maximizing expected somethings, nor max-min-ing, nor "gaming it," to reveal the preferred strategy. How can he find the

optimal course of action to recommend to his decision-maker?

The simple answer is that he probably can't. The same answer is also the beginning of wisdom in this business. There has been altogether too much obsession with optimizing on the part of operations researchers, and I include both grand optimizing and sub-optimizing. Most of our relations are so unpredictable that we do well to get the right sign and order of magnitude of first differentials. In most of our attempted optimizations we are kidding our customers or ourselves or both. If we can show our customer how to make a better decision than he would otherwise have made, we are doing well, and all that can reasonably be expected of us.

And this much we frequently can do. It is much easier to find a system or strategy that dominates (or nearly dominates) some other system -- say the one currently planned -- than it is to find a system that dominates all other systems. There is great scope here for the use of the a fortiori argument, for showing that even in the less favorable contingencies your system does as well or better. And if we can't find such a system we can frequently invent one. In fact, it has been our experience that ingenuity is frequently more profitably exercised in invention than in mere judging. If System A is good in some likely contingency 1, and very poor in another likely contingency 2, and System B vice versa, the really useful thing for the operations researcher to do isn't to make the perfect choice between A and B, but to invent or develop a System C that will perform reasonably well in both contingencies. If he can't do anything else, he can find a means of providing some insurance against the uncovered contingency, and calculate its cost.

Another thing that he can frequently do, especially in problems involving research and development, is to ascertain the critical uncertainties and recommend strategies to reduce them -- to buy information. If we don't know which of two dissimilar techniques for missile guidance will turn out to be better, the best recommendation is very likely to be: keep them both in development a while longer and choose between them when you have more and better information. Never mind the people who call you indecisive. You can prove that this kind of indecisiveness can save both money and time.

Of course you can't afford to try everything. There isn't enough budget. There aren't enough resources. You remember when we used to say "If you gave the military services everything they asked for they'd try to fortify the moon!" (We'll have to change that figure of speech.) Actually, that's why operations research and operations researchers are important. There'd be no problems for us if there were no constraints on resources. It is our job and opportunity to find, or invent, within the constraints, some better pattern of adjusting to an uncertain world than our betters would find if we weren't here; or some better way, taking costs and payoffs into account, to buy information to reduce the uncertainty.