

604211

604211

0

SUB-OPTIMIZATION CRITERIA
AND OPERATIONS RESEARCH

Roland N. McKean

P-586 ✓ *Bea*

22 April 1953

Approved for CTS records

211

COPY	<i>1</i>	OF	<i>1</i>
HARD COPY	\$	<i>1.00</i>	
MICROFICHE	\$	<i>0.50</i>	

DDC
RECEIVED
AUG 19 1964
DDC-IRA C

The RAND Corporation

1700 MAIN ST • SANTA MONICA • CALIFORNIA

SUB-OPTIMIZATION CRITERIA AND OPERATIONS RESEARCH*

Roland N. McKean

First of all I should like to thank Dr. McCloskey and the others who have arranged this opportunity for me to speak to you. In introducing my subject, I won't spend much time trying to define Operations Research. I'm sure that is unnecessary in the 24th lecture of this series. I think of Operations Research as simply the use of scientific method to help us choose among alternative means (including newly devised means) to our ends. These "means" may be alternative motions, factory layouts, operational procedures, objects of expenditure, allocations of a budget among those objects of expenditure, and so on.

Now in this process of choosing among alternative means to our ends, it is pretty obvious that we need to scan the ends themselves with a critical eye. This is a thesis that is emphasized a great deal by philosophers and theologians. They remind us frequently that our industrial economy and scientific developments may be extremely efficient in achieving certain ends but that these ends may be the "wrong" ones -- ends that are selected almost unconsciously or at least without sufficient critical thought. Philosophers stress that we must keep reviewing and questioning the pattern of ends implicit in our behavior.

A homely illustration that is often cited is that of a football player making a brilliant 90-yard run to the wrong goal. He may make a series of marvelous sub-optimizations and perform some highly efficient footwork, but

*

This address simply elaborates some of the points made by C. J. Hitch, in "Sub-optimization in Operations Problems", Journal of the Operations Research Society of America, May, 1953, pp. 87-99. I am indebted also to W. Capron, G. Cooper, and J. Hirshleifer of The RAND Corporation.

since he has been careless in selecting his goal, the total performance somehow doesn't add up to an efficient one.

Frank Knight, a well-known economist at the University of Chicago, often emphasized the importance of ends, and the danger of neglecting them, in our notions about efficiency or economy. Efficiency just doesn't mean anything except in terms of values or ends or objectives. When we talk about the efficiency of a steam engine, we are talking about useful output in relation to input. If we drop the word "useful", the term "efficiency" loses all sensible meaning: the law of the conservation of energy tells us that all inputs become outputs in one form or another so that without the notion of useful output, efficiency is always 100 per cent. Hence, we have to watch our step in defining useful output. The steam engine may be relatively inefficient in getting usable energy out of fuel, yet relatively efficient in making profits for a particular establishment. In brief, we have to keep our value-judgments or objectives constantly in mind.

Of course, the concept of ends or objectives is not just a simple listing of several desirable events or desirable outcomes. Probably the only sensible way we can visualize ends or objectives is as a preference surface or indifference surface relating a group's welfare, a firm's profit, or an individual's utility to various combinations of goods and events. For this reason, I will usually substitute for the term "ends" or "objectives" such terms as "utility function", "profit function", or "welfare function".

I have just been emphasizing the importance of ends or utility functions in connection with choosing among alternative courses of action. Criteria in operations research are important in much the same way. Criteria are in fact the practical counterparts of, or substitutes for, the functions that we would like to maximize in choosing among alternative courses of action.

In comparing alternative policies or operations, one would like to project, if it were possible, the total amount of utility, profits, or welfare under each of the various alternative policies. Then one would pick the policy that yielded the most utility, the most profits, or the most welfare (depending upon whether the "chooser" was an individual, a firm, or a government, respectively). In real problems, however, we have to look at some proximate criterion which serves, we hope, to represent or reflect what would be happening to utility, profits, or welfare. Hence, criteria -- in other words, what we look at in order to distinguish better from worse policies -- are not actually the utility or profit functions that we would like to maximize, but are proximate indicators.

Let me illustrate this distinction with an example. Suppose an industrial-machinery firm wanted to compare alternative methods of delivering parts to customers. Its researchers might adopt any one of a number of possibilities as a criterion to indicate the preferred delivery system. To mention a few of these possibilities, the criterion might be maximum expected profits, if they could be projected. Or, in lieu of this, it might be maximum revenues for a given cost or minimum cost of generating a given volume of revenue. It might be maximum volume of deliveries (in a specified time period) for a given cost, or minimum cost of achieving a specified type and number of deliveries. Thus, it wouldn't always be accurate to say that the actual criterion was the maximization of the profit function, though ideally this would be the proper criterion. The most accurate way to use the term, I think, is simply the way the dictionary uses it: a criterion means some rule or test by means of which we try to select the preferred situation or optimum course of action. To repeat, a criterion doesn't usually mean, in practice, the profit function or the utility function, or that magnitude

which we would in fact like to maximize. It is generally some compromise, some proximate indicator that is intended to lead us toward the maximization of whatever we would really like to maximize. To return to the dictionary definition, criteria are simply rules or tests by means of which we select preferred policies.

Now the decisive importance of criteria-selection, just like the importance of our choice of ends in general, ought to go without saying. Yet there is a great danger of forgetting, or at least neglecting, the significance of criteria-selection. If extreme care is not exercised in this part of the task, all the researcher's ingenuity and scientific tools may be wasted in deriving right answers to the wrong questions — which are sometimes diametrically wrong answers to the real questions.

Let me illustrate how wrong a decision-maker can be if he adopts an erroneous criterion with an example that you might say comes from municipal government. I heard of a motorcycle policeman once who patrolled a little stretch of road outside Chicago. It was a speed trap, and one motorist kept getting caught every few days by this cop. He got fed up with it, and had his car souped up so that it would do 120 miles an hour. Sure enough, the next day the cop started chasing him. As he began to close up on the speeder, the motorist suddenly stepped on the accelerator and disappeared. In the next scene, a number of people were gathered around a tree at the side of the road, inspecting a smashed motorcycle, and trying to help the unconscious policeman. They finally revived him and asked, "What happened?" The cop looked around and then answered, "I can't understand it. I was just about to catch that speeder and all of a sudden my motorcycle stopped. I got off to see what was the matter and that's the last thing I remember."

Now the criterion on which he based his choice among alternative actions was evidently the rate of change in the gap separating the two vehicles, or,

we might say, the ratio of the automobile's speed to the motorcycle's speed. It was a wrong criterion, and it led to an incorrect course of action.

But, in all seriousness, the selection of appropriate criteria is a crucial phase of any piece of operations research. This phase may not involve the use of scientific method -- it may not be an exciting endeavor -- but it is a part of the operations researcher's task. I realize that scientific analysis is supposed to be relatively independent of general value-judgments or objectives. In economics, we often analyze the consequences of alternative policies, and then say, "The voters must choose among these alternatives in the light of their own value judgments." Or maybe, "If your objectives are thus and such, you should adopt policy A, but if your objectives are so and so, you should adopt policy B." In consumers' research, also, scientific tools are used to describe the characteristics of alternative products, but the researcher does not presume to set down the consumer's utility function, and indicate the optimal use of the consumer's funds. Even in these types of analysis, of course, the researcher has to decide which are the consequences or characteristics in which the consumer or voter is interested. But in most operations research, an attempt is made to distinguish better from worse policies, or to pick out the optimal course of action. In this type of work, while the customer can help the researcher by describing his general objectives, the researcher usually must and should select the proximate criterion. To repeat, this is a part of his task. Nobody is going to hand appropriate criteria to him on a platter.

C. J. Hitch has summed up much of the above as follows: "The validity and therefore the usefulness of operations research depend upon the skill with which projects are designed and particularly upon the shrewdness with which criteria . . . are selected. The criterion problem has been relatively

neglected in operations research literature, and has apparently usually been 'solved' in practice by assuming the first plausible payoff function which springs to mind; or if several spring to mind, by trying all and compromising (or letting a commander compromise) among the results of alternative computations. The problem is much too important for such casual treatment. Calculating quantitative solutions using the wrong criteria is equivalent to answering the wrong questions. Unless operations research develops methods of evaluating criteria and choosing good ones, its quantitative methods may prove worse than useless to its clients in its new applications in government and industry.* Thus, what I have been trying to say is by no means original, but in my opinion it is worth considerable repetition.

So far I have just been discussing the importance of criteria-selection in general terms. Its importance is brought out in the paper by Hitch, from which the above quotation was taken, partly by means of specific examples of operations research, and partly by means of a look at the nature of sub-optimization. I will attempt to elaborate a little upon the last point -- the nature of the optimizations that operations research can provide.

One of the facts of life that makes criteria-selection so difficult is that we always deal with incomplete optimization and sub-optimization. It will make for greater clarity, I think, if these two concepts are distinguished. Incomplete optimization will be taken up first.

Full optimization would require the simultaneous consideration of all possible allocations of one's resources -- that is, all possible alternatives and all possible allocations of one's resources among those alternatives. It would require consideration of the possible impacts of all exogenous

*

C. J. Hitch, op. cit., p. 87.

events -- that is, all events not under the optimizer's control -- and expectations about their occurrence. Finally, it would require the maximization of the utility function of the optimizer (e.g., an individual, the managers of a firm, or the decision-makers in a government) subject to certain initial constraints.

Now any piece of operations research falls short of this requirement in all three respects. In other words, we are always dealing with incomplete optimization. The reasons that any piece of operations research falls short are probably obvious, but let me go over them briefly. First of all, such research considers only a few of the possible alternatives and usually only a few allocations of resources among these alternatives in an attempt to select an optimum policy. For example, suppose we return to the industrial-machinery firm that is trying to compare alternative methods of delivering parts to its customers. The company officials or researchers are likely to consider only a partial list of alternatives. Assume for a moment that they compare the results of using railroad shipments with the consequences of owning their own fleet of trucks. Obviously, if they ignore a third method -- e.g., the possibility of using air freight -- they may not reach the best answer. At first glance, the solution may appear to be simple: just make sure that all alternatives which are close substitutes are considered. The difficulty comes about because of the fact that alternatives are not always obvious substitutes. They are often dissimilar in physical appearance and often differ in their specific function. Indeed, to turn immediately to the extreme case, all of the various things which a firm or individual buys, or might buy, are alternative objects of expenditure which may contribute to the firm's or individual's general objective, even though their specific functions are as different as day and night. Bookkeepers and overhead cranes -- pianos and gundrops -- all are alternatives, substitutable in varying degrees,

competing for the purchasing agent's check.

Now it is obviously impossible to consider simultaneously the whole range of alternatives. Moreover, it is surprisingly easy to overlook a crucial course of action which, if considered explicitly, changes drastically the preferred allocation of funds among the various alternatives.

The second reason we are always dealing with incomplete optimization -- the first being that we can't possibly consider all of the relevant alternatives -- is that one can make only a few assumptions about events other than those controlled by the optimizer. Let me turn, once again, to the hypothetical industrial-machinery firm trying to compare alternative means of delivering parts to its customers. Now some action not explicitly considered might affect revenues or costs with railroad shipments differently from the way it would affect revenues or costs with the firm's own fleet of trucks or with air freight. For instance, management might decide to relocate the plant or build some additional plants. Operations might suddenly expand on the West Coast, or the local market might boom. Such events beyond the control of the optimizer -- decisions at higher administrative levels, or developments outside the firm -- can turn the researcher's answers upside down. Yet, even if he could quantify expectations about such occurrences, it would be impossible to take into account all possible assumptions about these contingencies. Again, it simply means that operations research can be no more than incomplete optimization.

A third reason, a very important one, that we can attempt only incomplete optimization, is that we must inevitably use an imperfect criterion. As stated earlier, what one would like to examine, under the alternative operations, is the individual's utility, if the optimizer is an individual consumer -- or the firm's stream of profits, if it is a firm. We can't do that, of course. Profit maximization, by the way, is not much simpler

conceptually than utility maximization. We often say that what really matters to a firm is the maximization of profits, but this doesn't make sense in an uncertain world. If the managers consider several alternative courses of action, maybe they can anticipate the amount of expected profits in various future periods under each alternative policy. Maybe they can say something about the variance associated with the expected amounts. In other words, maybe they can say to themselves, "Under Policy 1, our expected profits will be X_1 in the first period, X_2 in the next period, and so on. Under Policy 2, our expected profits will be Y_1 (less than X_1) in the first period, Y_2 (greater than X_2) in the second period, and so on. Under Policy 3, our expected profits are greater in every period than they would be under the other policies, but there is a greater risk, on the other hand, of suffering huge losses." What does it mean to maximize profits under those circumstances? All it can mean is the maximization of some function of the profits expected in various periods and with various "probabilities". This function depends upon management's attitude toward risk and toward profits now versus profits five years hence. We don't know this function any more than we know an individual's or a government's utility function.

But even if we had a clear notion of these functions, we would have to settle for a more manageable, albeit a less appropriate, criterion instead of looking at the ultimate impact on profits or utility or social welfare. For example, as the officials in our hypothetical firm look at various delivery systems, can they really translate the results under each delivery system into effects on expected sales, expected revenues, expected costs, and,

*

See A. Alchian, "Uncertainty, Evolution, and Economic Theory", The Journal of Political Economy, XVIII (June, 1950), 211-21, or S. Enke, "On Maximizing Profits", American Economic Review, XLI (September, 1951), 566-78. This same question will harass the researcher in the selection of the "best" stream of costs, gains, profits, utility, welfare, or anything else.

ultimately, expected profits? Probably not — at least with any confidence. So they settle for something a little less appropriate -- but something they can handle. They may try, for example, to select the delivery system that minimizes the costs of achieving, say, 48-hour delivery service. They may try to select the method that minimizes the cost of delivering parts, regardless of the by-product effects on the cost of delivering other items. Inevitably, the criterion will be imperfect. It will not reflect the full range of effects on revenues, costs, and total profits. So, for still a third reason, operations research can yield no more than incomplete optimizations.

Thus, the methods of science will not magically or automatically reveal optimal choices. It is manifestly impossible to demonstrate that incomplete optimization, which is what we always do, will in general lead toward the maximization of whatever we would like to maximize. Some incomplete optimizations are better than others; some may give worse results than flipping coins. It will always be necessary to use great care and good judgment in picking out the partial optimizations that look promising and to set up appropriate criteria in terms of which to judge the alternative policies or courses of action.

I turn now to a special kind of incomplete optimization which has been termed "sub-optimization". An essential aspect of governmental or business operations is that decision-making is decentralized in varying degrees. Thus, different administrative levels allocate the resources at their disposal among the uses under their supervision. The government provides one of the best illustrations. At relatively high levels the Federal budget is allocated among Departments, such as Interior, Defense, and Agriculture; among the independent agencies; and among broad programs. The allocations of

funds within Departments and programs, among different projects, and amongst different operations within each project, are accomplished by administrators at "lower levels" — in other words, by different people who face narrower problems.

In business this is true, too. In large corporations, authority has to be delegated. It is impossible for top management to make all the decisions and optimize over the whole list of alternative operations or courses of action simultaneously. So the corporation is divided into divisions or departments. The broad decisions are made at the top level, but each department and division must choose among narrower alternatives, and so on down the line to foremen and lower-level administrators who must choose among alternative schedules of operations, employees, motions, and so on.

Even an individual may handle his affairs in this way. He may first allocate his income or budget among the big categories, such as food, clothing, and housing. He may then delegate to himself in another capacity, so to speak, the task of selecting the individual suits, pairs of shoes, and cuts of meat that fall under these broad headings. He may adhere for the most part to the original allocation among the broad budget categories. In other words, he may allocate his income among these categories at a "higher level" and then perform narrower day-to-day optimizations at a "lower level". It would be impossible, when making individual purchases, to compare them with all alternatives and re-examine each day one's allocation among all major objects of expenditure.

This "lower level" optimization is what I will refer to here by the term "sub-optimization". In other words, sub-optimization, as used here, means the process of choosing among a relatively narrow list of operations or resource allocations by an administrative level other than the highest.

Careful selection of criteria in sub-optimization is difficult, and it is crucial. For unless the criteria are selected with great skill, operations research may increase a department's efficiency in terms of its test, yet decrease the firm's efficiency in terms of its ability to make profits. Since our research always pertains to sub-optimizations, this complication in criteria-selection is always present.

One reason that sub-optimizations present these difficulties is that it may be computationally impossible for department heads to compare the alternatives facing them in terms of the criterion that would be appropriate for the entire firm or government. As the firm grows, the difficulties of coordination increase, and genuine confusion develops about objectives at departmental and lower levels. The chances of adopting inferior criteria multiply. If the sales department fights to maximize the ratio of sales per dollar of selling expense, or the production department tries to minimize cost per unit of output, or the foreman tries to minimize the number of "seconds" or "rejects" per month, it may cause the firm as a whole to sacrifice profits. In government, where it is usually impossible to measure the counterpart of a firm's profits, sub-optimization criteria can get still further out of hand. The criterion of an office's efficiency may be the number of licenses granted per dollar spent or the number of letters written per typist. Certain types of equipment may be selected because they yield maximum expected life or speed per piece of equipment or maximum accomplishment per man employed. Such criteria, particularly ratios,* may not lead to increased efficiency in terms of what the government really wants to do.

But even if we avoid downright errors about the nature of appropriate criteria, sub-optimization poses obstacles to other aspects of criteria-selection.

* For a more thorough discussion of ratios' shortcomings, see Hitch, op. cit., p. 94.

In general terms, the appropriate criterion for the firm, government, or individual seems to be the maximization of gains minus costs (not the ratio of gains to costs). Of course, this statement doesn't really tell us much. In any particular piece of operational analysis, the big questions still remain. One of the big questions is what should in fact be counted as the costs of the alternative operations, programs, or policies? The costs of one program ought to measure the gains that might have been obtained by using the same resources in other ways. In brief, costs are what you give up in order to have this program. For instance, the costs of operating your own fleet of trucks are whatever you have to give up in order to do so. If you have to buy them, the cost is the dollar outlay over the relevant period (minus salvage value at the end of the period). In many proposed operations or programs, however, large portions of the equipment may already be on hand. If an item that would be used in an operation is in stock, what should its cost be? Presumably not the original cost. What is relevant at this point is its value in alternative uses. Suppose the hypothetical firm that was comparing alternative methods of delivering parts had some trucks standing around. What would be the cost of using its own fleet of trucks to deliver parts? If the only alternative use of the trucks on hand was to sell them, then their salvage value would be the cost that should be charged for them in costing this particular delivery system. If, however, they could be used to advantage in another department, their value in this alternative use might be greater than salvage value. In sub-optimizations, it is easy to overlook the worth of equipment, personnel, or any resources in other departments. Actually, this is not usually an important consideration in the private firm, for free markets provide so many yardsticks to help guide each departmental "sub-optimizer". But it may be an important consideration in government,

particularly in the costing of alternative military operations.

Sub-optimization -- the fact that lower administrative levels set about their operational analysis more or less independently -- aggravates another problem in costing any operation. This compartmentalization of decision-making increases our tendency to ignore the indirect effect of one operation on costs of other operations. Assume that the delivery department wanted to choose the cheapest way of achieving a given objective (e.g., 48-hour delivery service to all customers). They might consider only costs to that department, yet the delivery systems might have various impacts on the costs of other operations. The use of its own trucks might decrease production costs by permitting more flexibility in choosing new plant locations, or the method considered might increase costs (e.g., for transportation) in other departments. For example, the use of trucks for some deliveries or some pick-ups might necessitate less-than-carload shipments of other items. To the individual department in charge of the operations, the impacts on its costs are felt directly and are more likely to be considered than costs felt by other departments. To the firm as a whole, of course, the criterion for choosing among operations would involve the net impact of each operation on total costs. Clearly it is important for sub-optimizing administrators to make the cost-aspect of their criteria consistent with the cost-aspect of higher-level criteria.

As Hitch pointed out, this is analogous to a subject that is frequently discussed in economics, the discrepancy between private costs and social costs. That is, the difference between the costs to the individual department and the costs to the whole firm is analogous to the discrepancy between the costs of some operation to a firm and the costs to society. One of the traditional examples is the discrepancy between the cost of an industrial plant's

use of soft coal as felt by the firm, on the one hand, and as felt by the entire economy, on the other. The cost to the entire group includes not only the costs to the firm but also the costs to "innocent bystanders" -- e.g., the increase in laundry bills and possibly in doctors' bills.

Economists usually say that resources would be better allocated among alternative uses if the firm's criterion and our over-all criterion could be brought into closer harmony by making firms feel the major indirect costs of their operations.

So much for the analogy. In any event, it seems clear that sub-optimizations within firms and particularly within governments often neglect this aspect of an appropriate criterion -- the costs imposed on other operations that are not being considered directly.

Let's turn to the other general aspect of appropriate criteria -- the gains from the operations being considered. What should be included in the calculation of gains? Often in operations research we "set" the gains -- i.e., specify a fixed objective -- and try to determine the cheapest way to achieve those particular gains. Now this is by no means the ideal criterion. Most firms wouldn't be content merely to discover the cheapest way to achieve a given volume of revenue. Nonetheless, this sort of criterion is often satisfactory, and it will serve here to bring out the problems of sub-optimization.

The first problem that is aggravated by sub-optimization is the selection of a sensible fixed objective. It is difficult, and dangerous, to narrow the proximate objective down to a single variable. The hypothetical firm that I have used to illustrate various points was comparing alternative methods of delivering machine parts. If their operations researcher tried to set up a given objective, he could not really use such a goal as 48-hour

service to all customers. In addition to speed of delivery, the department would be concerned also with reliability of delivery and with the condition of the parts upon arrival. The operations researcher, if he could not do better than choose some fixed proximate objective, would have to ask, "What objective makes sense from the firm's standpoint, in view of the probable effects of delivery upon sales, revenue, and ultimately profits?"

Another example may illustrate more vividly the necessity of looking at more than one variable and of being careful in setting up a given multi-dimensional goal. Suppose petroleum companies were interested in alternative ways of storing petroleum products. They might run a series of experiments to compare the consequences of storing products in various ways. They might specify an objective and choose the storage method that minimized the cost of achieving that objective. But what kind of an objective? A specified gum content after two months' storage? Or after six months storage? Should the objective include a measure of sedimentation or color? A certain ash content? A petroleum expert could reel off a dozen or more variables which might be included as dimensions in the specified goal. The sub-optimizer has to do his utmost to select this aspect of the criterion in the light of higher-level criteria -- or, in other words, from the firm's rather than a narrower viewpoint.

In some cases like the above, the quantitative answers that the operations researcher can provide may be only "consumers research", pointing out the cheapest ways of achieving several possible objectives or spelling out the gains (not a given objective) and costs of the alternative policies. This is what consumers research usually does for the consumer. It points out that TV Set No. 1 has a certain pattern of characteristics, while Sets No. 2 and No. 3 have certain other characteristics. The consumer then makes his own

selection. Of course, the operations researcher can, and usually should, offer his recommendation to his customer. But the researcher's recommendation and the client's ultimate decision will both be better if the main uncertainties are recognized. And if neither the client nor the researcher can pin down a unique specified objective, the cheapest way to achieve an assumed objective may be an answer that has spurious definiteness. As far as definiteness is concerned, our sub-optimizations are always incomplete, regardless of our ability to specify a unique objective. In almost no case does operations research provide "the" answer; it is always a little like consumers' research, providing quantitative information which must be supplemented by the exercise of judgment in order to yield recommendations.

To sum up this last point, it is not a simple matter to perceive a measure of gains for lower administrative levels that will harmonize nicely with objectives at higher levels. The necessity of sub-optimizing helps to fill this phase of criteria-selection with problems.

But sub-optimization still further compounds the difficulties of measuring the gains from alternative operations. In measuring gains as well as costs, we need to take into account the indirect impact on operations other than the ones explicitly being compared. The operations under consideration may provide benefits to other departments. If our hypothetical firm delivered parts with its own fleet of trucks rather than via air freight, it could presumably use the sides of the trucks for advertising. These gains would occur in a different department. They might conceivably be substantial. The operations researcher would have to take the firm's rather than some narrower point of view, and attempt to take such indirect gains into account. If the researcher had set up a given objective and was trying to determine the cheapest means of achieving it, such bonus or "external" gains would

have to be evaluated and deducted from the cost of the operation that yielded the bonus gains.

Again, there is an analogous problem in economics that has caused a good deal of discussion. The discrepancy between gains to a department and gains to the entire firm or government is analogous to the discrepancy between private gains and social gains. That is, the gains to society of a firm's operation (e.g., a training program for machinists) may exceed the gains to the firm itself. This analogy serves perhaps to re-emphasize the point -- that in any sub-optimization, the operations analyst needs to make a special effort to check for possible gains to other operations as a result of the ones under consideration.

Everything that I have said can be summed up in the following propositions. Criteria-selection is an inescapable and a crucial part of the operations researcher's task. The fact that we always sub-optimize helps lure us toward erroneous criteria. The fact that we sub-optimize means that we must always keep a weather eye out for (1) the value of resources if they were used in other departments, (2) the impact of any operation on the costs and gains from other operations, and (3) the consistency of the criterion (embracing the measure of gains and the measure of costs) with the aim of the firm (or government) as a whole. In other words, the sub-optimizer should adopt for his analysis the same best criterion that top-level executives could adopt if they had to examine that particular problem.

The compatibility of lower-level with higher-level criteria is particularly significant in another operations research problem. In some situations, operational analysis may recommend more extensive use of decentralized decisions guided by prices. Such arrangements may be helpful occasionally within huge firms, but more frequently within institutions or governmental units.

For example, the Defense Department is using stock and industrial funds to manage numerous activities as though they were private manufacturing or retailing establishments. Instead of handing out clothing, printing services, or Military Sea Transport Service for free whenever requisitioned, the Navy sells these items to various military units. This forces the users, in making their decisions, to take into account the fact that these items cost something -- i.e., that they have alternative uses.

Now, in what situations will this arrangement function effectively? Operations research could presumably help answer this question. It is apparent, however, that the answer hinges on the criterion problem in sub-optimizations. Wherever this arrangement induces lower-level sub-optimizers to adopt criteria that harmonize with top-level criteria, it seems promising. If the prices charged, and the purchasing power allocated among the users, are such as to make criteria more nearly compatible, this institutional change would probably bring about genuine economies. If the situation or the method of pricing is not conducive to the consistency of lower-level criteria with higher-level aims, then other institutional arrangements are likely to be superior. For instance, if GI's buy their clothing from an allowance, they will be induced to care for their clothes in a fashion that corresponds closely with the aims of the Defense Department. On the other hand, if they bought their rifles, hand grenades, and other arms from the same allowance, they might be induced to make choices that harmonized rather poorly with top-level objectives.

Thus, a constant awareness of higher-level criteria is essential to useful lower-level sub-optimization. Discernment of top-level criteria may pay off also in another type of effort that is closely related to operations research. We hear a great deal about "performance budgets" these days --

partly for industry though chiefly for governments and institutions. Such budgets would not present management with optimizations or even recommendations; they would summarize the proposed expenditures on major functions or objectives and attempt to trace out quantitatively what those outlays would buy. I sometimes call this aggregative consumers' research. But it is closely akin to the conventional description of operations research. Scientific techniques help us to trace out the quantitative implications of the amounts spent on major functions. Also -- and very important from the standpoint of this paper -- the criterion problem confronts us again. Into what major functions or objectives should the proposed expenditures be split? What sort of quantitative implications -- i.e., performance units or program units -- would be most useful to Congress or top management? In trying to answer these questions, one must wrestle with the criterion problem: what ought to count as gains and costs under the alternative programs? In a performance budget, the dollar amounts would reflect projected costs; the functional breakdown and quantitative indicators would provide clues to performance or gains. Only with considerable insight into higher-level criteria can helpful quantitative clues and functional breakdowns be provided.

To conclude -- we have to live with decision-making by sub-optimization. In this process, careful attention to criteria -- the tests that are used to pick out preferred actions -- is imperative. And this means that in all our research -- whether we make lower-level sub-optimizations, do higher-level consumers' research, or recommend changes in the institutional arrangements for sub-optimization -- the researcher and client must keep probing into higher-level objectives and criteria.