COMPRESSED POLICY ANALYSIS

M. Fiering

March 1974 WP-74-10

5.4 Compressed Policy Analysis

5.4.1 Objective

(a) Some Definitions

The objective of this portion of the analysis is to seek approximations to an optimal strategy of forest management, and thereafter systematically to evaluate any such approximation in the hope that one or more might be sufficiently reliable to preclude the enormous computing effort subtended by those more comprehensive mathematical formalisms (linear and dynamic programming) which are also investigated in this study. Compressed Policy Analysis (CPA), unlike its more formal counterparts, does not identify an optimal solution, but rather provides a mechanism for rapid and fluent examination of alternative solutions which are generated exogeneously (in accordance with some systematic or randomized sampling procedure), and then, by a comprehensive display of economic criteria and other relevant performace factors, suggests to the decision-maker(s) which alternative to adopt.

A solution is a set of decisions germane to forest management. Certain options are available; these include a variety of timber cutting and harvesting patterns, rates of application of insecticide, and enhancement of wildlife and recreation facilities. The thrust of CPA is the identification of several policies which are deemed, *a priori* to be politically, socially and institutionally feasible, and the evaluation of these policies in ways which take explicit account of the multi-dimensional and highly variegated outputs of the forest ecosystem. Among the requirements of a feasible policy is that its action display a reasonable measure of spatial and temporal homogeneity. This tends to keep costs down, even though the policy consequences might be inhomogeneity in system response. Thus the richness of the potential policy space is compromized by practicality, and the full palette of a mathematically intact search for the optimal policy might prove too ambitious.

It is therefore proposed that CPA be applied to a relatively small set of policy options, including that one in current use, on the assumption that this set will be sufficiently fertile to identify some policies which are clearly inappropriate, some which merit further detailed consideration, and some to which system performance is largely indifferent. The point of the exercise is not to identify a sharply-defined optimum which might be extremely sensitive to unanticipated climatic or ecological perturbations but to sharpen the focus of subsequent debate by excising a small number of candidates for continuing analysis, and thereby to advance a general methodology for decision-making in an environmental context.

(b) A Formalism for Decision-Making

It would be ideal if mathematical programming could routinely be used to solve for the optimal policy under a variety of assumptions and conditions pertaining to our forest ecosystem model, but it should be recalled that policy evaluation proceeds within multi-dimensional space and this severly limits the applicability of techniques for direct identification of optimal policy. The multidimensionality of system outputs, undertainty about which outputs to include and how to weight them, and conflicts concerning the priorities expressed by the several claimants on the resource all conspire to make mathematical programming an unlikely tool for identification of the optimal policy in this forest management problem. Three alternative modes of analysis are suggested in this section;

all of them together are lumped under the rubric of Compressed Policy Analysis, and all suffer several disadvantages and imprecisions. But, as in virtually every real problem of policy analysis, there is no unique solution or method of analysis which clearly dominates the decision-making process; it is through the conjunctive use of exact and approximate solutions, computationally simple or exotic, deterministic or stochastic, descriptive or prescriptive, that grudging progress is made.

(c) Sampling in Policy Space

The basic tool for evaluation of policy options is simulation of the budworm-forest ecosystem. Initial applications used a short trace of meteorological inputs, with the model was applied to a single plot and with no spatial linkages to simulate pest dispersal through the entire region. Our work generalizes the program to accommodate dispersal over all 265 plots,

Simulation runs were made, each signed to test an alternative candidate for policy implementation. The candidates were developed after consultation with ecologist civil servants, representatives of recreation and wildlife groups, industrial proponents, and others whom we could identify as having a vested interest in management of the forest ecosystem. Due to the fact that this study identifies methodology rather than definitive conclusions, we did not pursue an extensive program of sampling in the space of policy options; restrictions on time and computing budget made this infeasible. Instead, we are concerned primarily with the exposition of a methodology for decision-making, so we present here a highly abbreviated examination of policy options. We did not use systematic or random sampling techniques for identification of policy options; for completeness of exposition these are described below. But even a limited analysis can be extremely useful if it turns out that system response, in terms of the output variables critical to the decision makers, seems to be relatively flat. That is, if it appears that response is not highly sensitive to a wide range of reasonable policy options, we might begin to appreciate, even from a superficial analysis, that it is unnecessary to undertake a very large random or systematic sample of policy options. It might turn out that there is enough buffering, enough natural resilience or persistence, in the system to conclude, or at least strongly to suggest, that the major issues are those of political acceptability rather than sensitivity to small changes in decision variables.

It is reasonable to ask how many policies or potential solutions should be investigated to be sufficiently certain that the sample from which our solution is drawn is big enough. Of course, in solution by mathematical programming, this question does not arise because the most commonly used techniques generate the optimal solution. But mathematical programming is not likely to be able to embrace the number of variables required for our forest ecosystem performance index; thus the generated policies are guidelines to the selection of a few policy options which can be further tested and refined by simulation. It is our intent that these few promising policies (or decisions) should form the basis of a more penetrating investigation which would lead ultimately to a final decision.

Techniques other than mathematical programming are available to identify candidates for simulation, and some of these are particularly powerful. For example, if the number of decision variables is small, and if each can be divided into a small number of alternatives, then it is reasonable systematically to examine all the intersections or potential decisions in multi-dimensional space, to evaluate each, and to pick the most promising few for further investigation. But it is in the nature of ecological systems that many dimensions are required, and it is unreasonable to divide all the decision variables into a small number of steps, so an exhaustive search for potential decisions cannot generally be undertaken. A particularly powerful tool under these circumstances is the use of random sampling techniques to develop trial solutions which can be improved by steepest ascent or "hill climbing" techniques now routinely used in applied mathematics. For example, we know that if a random sample of size n is taken, where each of the n points is another decision vector, then the probability is $1 - (1-\theta)^n$ that the best of all n trials lies in the upper θ-fractile of all possible results. This simple but powerful result is independent of the dimensionality of the decision and of the functional form of the distribution of any of the system outcomes. It requires only that the outcomes be represented on a continuum in multi-dimensional space, a condition which might sometimes be difficult to guarantee because of the potential lumpiness of system response. But experience with many resource investigations suggests that we can virtually always find reasonable and feasible policies (or solutions) which closely approximate the requirement that all outputs be defined on a continuum.

If we draw a random sample of size 30 and inquire about the probability that the best of these lies in the upper 10% of all possible results, we determine that the probability is 0.957. It should be emphasized that defining a point to lie in the upper θ-fractile is different than asserting that a point lies within θ of the true optimum. We make no statement here about the quantitative difference between the best of an independent random sample of outcomes and the true optimum; we define only the probability that a particular output lies within any given fractile. This probability, 0.957, is independent of the dimensionality of the decision vector.

Moreover, if the few best results of the random cast are systematically improved by various hill climbing procedures to promote them from their random positions in the decision space to a local optimum (from which all small changes make the output worse), we reside (symbolically) on a set of local mountain peaks from which all directions are down. From an operational point of view this is tantamount to saying that we have a new set of local optima, the best of which is at least as good as the best of the random draw because all movements were necessarily uphill (in the direction of increasing value of system output). There is no general theory which describes how to calculate the confidence and tolerance limits on the best of the local optima because such a result would depend on the nature of the response surface and on statements about higher derivatives.

The decision as to whether n, the size of the initial random sample, is large enough depends on the cost of drawing additional samples (that is, on the cost of computing) and on the fertility of our imagination (because it is required that these policy options be feasible and it is oftentimes difficult to generate feasible random combinations of decision variables). Thus one should not be misled by the apparent simplicity and elegance of random sampling techniques but should recognize that the potentially high cost of identifying a random feasible candidate may make the procedure unattractive. Moreover, it is obvious that the probability of identifying a feasible random solution becomes painfully acute as the number of dimensions, and possible interactions among decision variables, increases. All this has been investigated by many theorists, and the results are neither conclusive nor satisfactory; the best that can be done here is to use currently available techniques, modified by the best advice we can obtain from consultants and practitioners.