

A CASE STUDY OF
FOREST ECOSYSTEM/PEST MANAGEMENT

C.S. Holling, G.B. Dantzig, C. Baskerville,
D.D. Jones and W.C. Clark

June 1975

WP-75-60

In Press: Proceedings International
Canadian Conference on
Applied Systems Analysis,
1975.

Working Papers are not intended for
distribution outside of IIASA, and
are solely for discussion and infor-
mation purposes. The views expressed
are those of the author, and do not
necessarily reflect those of IIASA.

INTRODUCTION

The boreal forests of North America have, for centuries, experienced periodic outbreaks of a defoliating insect called the Spruce Budworm. In any one outbreak cycle a major proportion of the mature softwood forest in effected areas can die, with major consequences to the economy and employment of regions like New Brunswick, which are highly dependent on the forest industry. An extensive insecticide spraying programme initiated in New Brunswick in 1951 has succeeded in minimizing tree mortality, but at the price of maintaining incipient outbreak conditions over an area considerably more extensive than in the past. The present management approach is, therefore, particularly sensitive to unexpected shifts in economic, social and regulatory constraints, and to unanticipated behaviour of the forest ecosystem.

Most major environmental problems in the world today are characterized by similar basic ingredients: high variability in space and time, large scale, and a troubled management history. Because of their enormous complexity there has been little concerted effort to apply systems analysis techniques to the coordinated development of effective descriptions of, and prescriptions for, such problems. The Budworm-forest system seemed to present an admirable focus for a case study with two objectives. The first, of course, was to attempt to develop sets of alternate policies appropriate for the specific problem. But the more significant purpose was to see just how far we could stretch the state of the art capabilities in ecology, modelling, optimization, policy design and evaluation to apply them to complex ecosystem management problems.

Three principal issues in any resource environmental problem challenge existing techniques. The resources that provide the food, fibre and recreational opportunities for society are integral parts of ecosystems characterized by complex interrelationships of many species among each other and with the land, water and climate in which they live. The interactions of these systems are highly non-linear and have a significant spatial component. Events in any one point in space, just as at any moment of time, can affect events at other points in space and time. The resulting high order of dimensionality becomes all the more significant as these ecological systems couple with complex social and economic ones.

The second prime challenge is that we have only partial knowledge of the variables and relationships governing the systems. A large body of theoretical and experimental analysis and data has led to an identification of the general form and kind of functional relations existing between organisms. But only occasionally is there a rich body of data specific to any one situation. To develop an analysis which implicitly or explicitly presumes sufficient knowledge is therefore to guarantee management policies that become more the source of the problem than the source of the solution. In a particularly challenging way present ecological management situations require concepts and techniques which cope creatively with the uncertainties and unknowns that in fact pervade most of our major social, economic and environmental problems.

The third and final challenge reflects the previous two: How can we design policies that achieve specific social objectives and yet are still "robust"? Policies which, once set in play, produce intelligently linked ecological, social and economic systems that can absorb the unexpected events and unknowns that will inevitably appear. These "unexpecteds" might be the one in a thousand year drought that perversely occurs this year; the appearance or disappearance of key species, the

emergence of new economic and regulatory constraints or the shift of societal objectives. We must learn to design in a way which shifts our emphasis away from minimizing the probability of failure, towards minimizing the cost of those failures which will inevitably occur.

The Descriptive Analysis

The descriptive analysis of the budworm/forest system aimed to produce a well validated simulation model that could be used as a laboratory world to aid in the design and evaluation of alternate policies. The key requirement of that laboratory world is that it capture the essential qualitative behaviour of the budworm forest ecosystem in both space and time. Extensive data concerning forest-pest and economic interrelations had been collected over the past 30 years by Environment Canada as one of the earliest interdisciplinary efforts in the field of renewable resource management. There are many missing elements, but this is an inevitability rather than a draw-back. If systems analysis is to be applied successfully to the management of ecological systems, it must be able to cope with unknowns.

The essential qualitative behaviour in time has been identified through an analysis of tree ring studies. Four outbreaks have been detected since 1770, (Fig. 1), each lasting 7 to 16 years, with a 34 to 72 year period between them. During the inter-outbreak periods the budworm is present in barely detectable densities which, when appropriate conditions occur, can increase explosively over four orders of magnitude during a 3 to 4 year period.

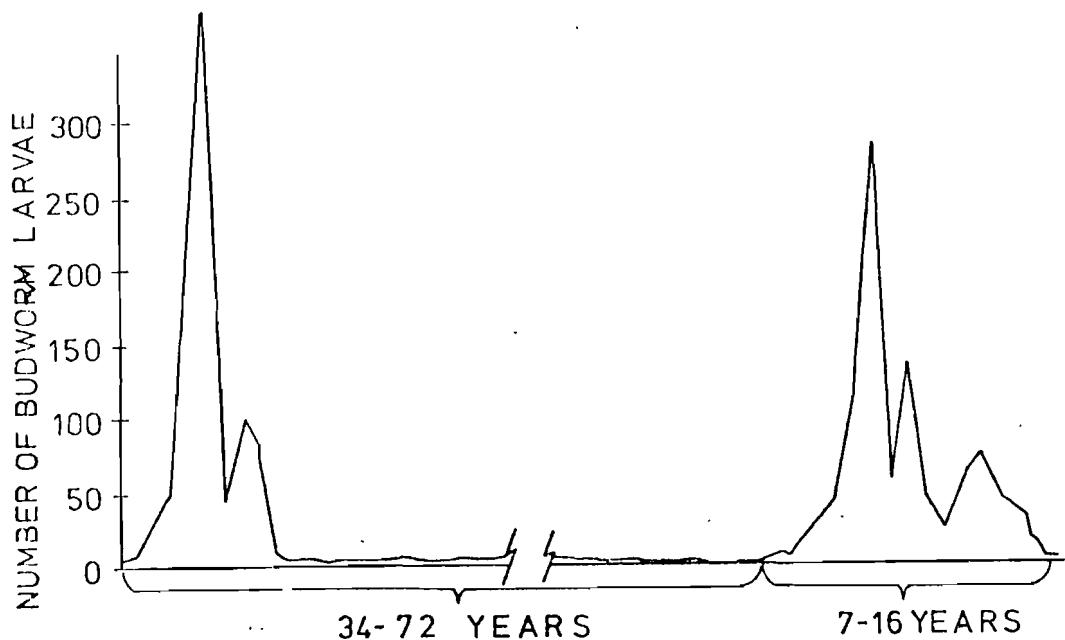


FIGURE 1.: The pattern in time. Representative historical pattern of Spruce Budworm outbreak. There have been four major outbreaks since 1770.

TABLE I

The State Variables Emerging From The Bounding
of the Problem

IDEAL NUMBER OF STATE VARIABLES

In one subregion	1
Birch	1
Spruce by age	30
Balsam by age	70
Budworm	1
Natural enemies	1
Weather	1) 1) retains memory
Tree stress	1)
Foliage new	1
Foliage old	1
Number of state variables per subregion	107
Total number of state variables in all 265 subregions	107 x 265 = <u>28,355</u>

The distinctive pattern in time is paralleled by one in space. The historical outbreaks typically initiated in one to three or four local areas of Eastern Canada and from those centres spread and contaminate to progressively larger areas. Collapse of the outbreaks occurred in the original centres of infestation in conjunction with mortality of the trees and similarly spread to the areas infested at

}

later times. The resulting high degree of spatial heterogeneity in the forest age and species composition is closely coupled to the "contamination" feature caused by the high dispersal properties of this insect.

The essential first step in the dynamic description of this system is a parsimonious bounding of the problem in terms of prime variables, space and time. The process of bounding the problem from the very start of the analysis is a key activity. Everything else in the analysis flows from these decisions and they profoundly influence the final form and relevance of the policies. The key requirement in bounding the problem in space, time and variables is to ruthlessly simplify while still retaining the essential properties of behaviour and needs for management.

Bounding Time

Because of the pattern of outbreaks shown in Figure 1, the minimum time horizon required is that which can contain two outbreaks - that is 150 to 200 years. In order to capture the dynamics of this system it is essential to have a time resolution of one year with seasonal events implicitly represented.

Bounding in Space

As in many pest species, the budworm disperses over long distances. The modal distance of dispersal is about 50 miles from one site, but dispersal distances of several hundred miles have also been recorded. It was thought essential to have a minimum total area based on at least twice this modal distance, leading to a minimum modelled region of 14,000 to 15,000 square miles. The area chosen in this study

was a 17,000 square mile area containing much of the Province of New Brunswick (Figure 2). But even events in this size of area are profoundly affected by contagion from outside it. It was therefore necessary to add a buffer zone of approximately 75 miles width around the area in order to compensate for edge effects. The behaviour of this system is as highly heterogenous in space as it is in time, and because of the contagion problem spatial disaggregation is essential. There is high variation in the spatial distribution of the primary tree species, of harvesting activities and of recreational potential, in part as a consequence of the historical interplay between the forest and the budworm. The 50 mile modal dispersal distance also suggests a minimum resolution of about one-fifth to about one-tenth of that distance. Hence the overall area is divided into 265 distinct 6 by 9 mile subregions (Figure 3).

Bounding Variables

An ecosystem of this extent has many thousands of species and potential variables. Our understanding of the dominant budworm-forest dynamics is sufficiently detailed however, that the system's relevant behaviour can be captured through the interrelations among five species, each of which represents a key role in determining the major dynamics of the forest ecosystem and its resulting diversity. These key variables are summarized in Figure 4.

The principal tree species are birch, spruce and balsam. In the absence of budworm and its associated natural enemies, balsam tends to out-compete spruce and birch, and so would tend to produce a monoculture of balsam. Budworm, however, shifts that competitive edge since balsam is most susceptible to damage, spruce less so, and birch not at all. Thus there is a dynamic rhythm, with balsam having the

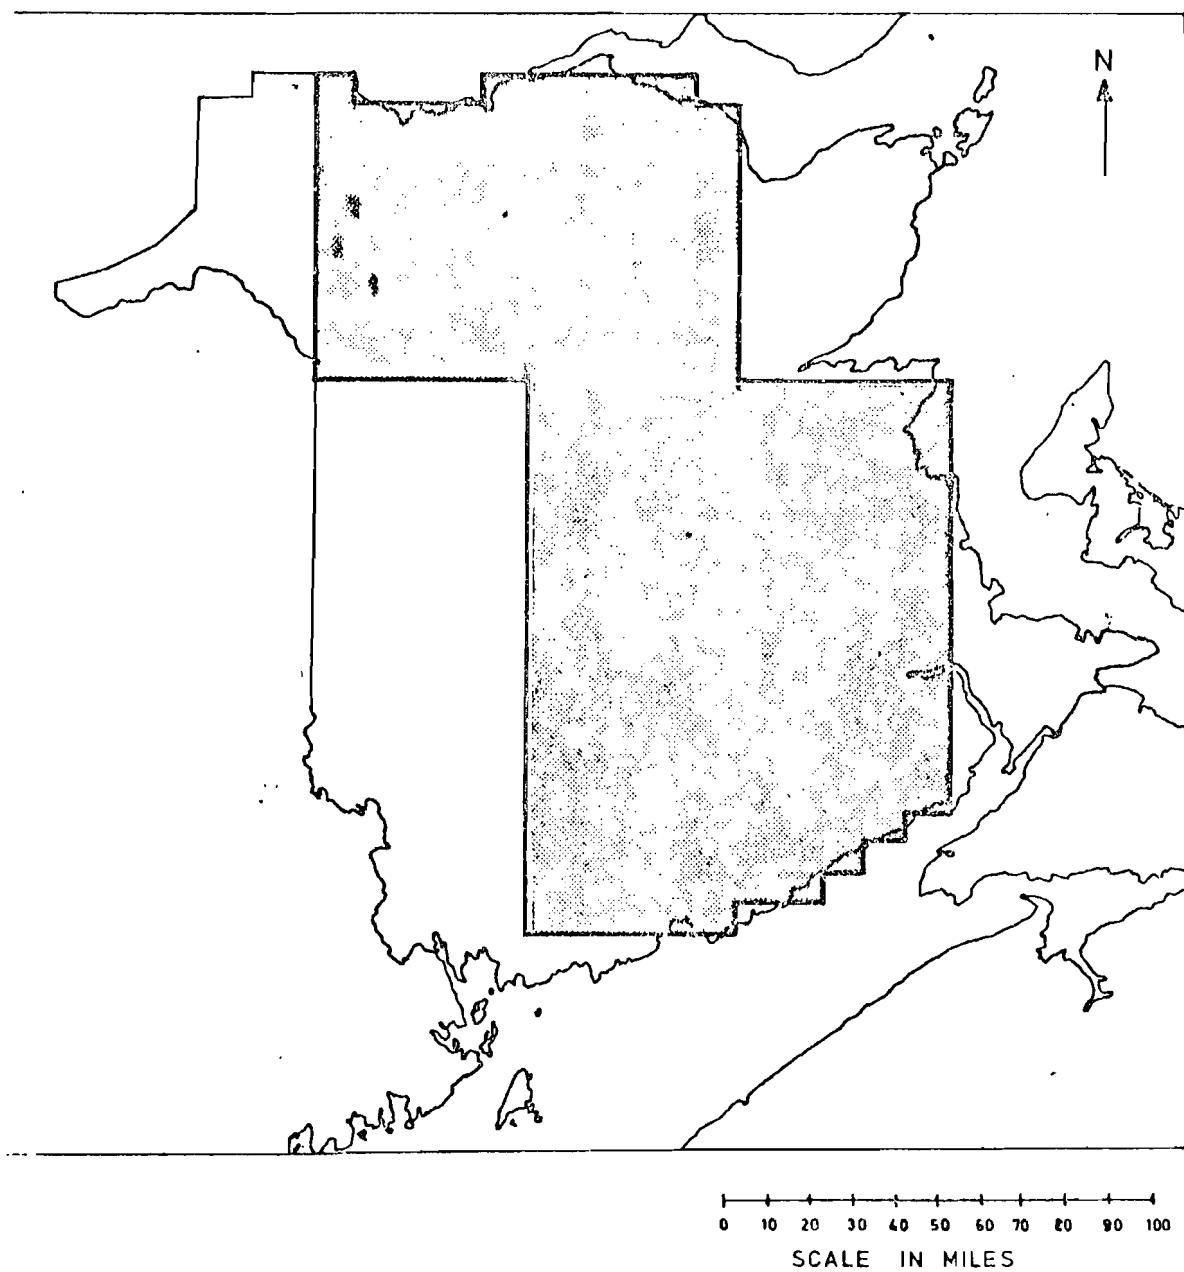


FIGURE 2.: Study area within the Province of New Brunswick used in the current study. The hatched area includes the primary forested regions of New Brunswick.

	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15					
01	1				2	3	4	5	6							2				
02	7	8	9	10	11	12	13	14	15	16						3				
03	17	18	19	20	21	22	23	24	25	26						4				
04	27	28	29	30	31	32	33	34	35	36						5				
05	37	38	39	40	41	42	43	44	45	46						6				
06	47	48	49	50	51	52	53	54	55	56						7				
07	57	58	59	60	61	62	63	64	65	66						8				
08	67	68	69	70	71	72	73	74	75	76						9				
09	77	78	79	80	81	82	83	84	85	86						10				
10	87	88	89	90	91	92	93	94	95	96						11				
11						97	98	99	100	101	102	103	104	105	106	12				
12						107	108	109	110	111	112	113	114	115	116	13				
13						117	118	119	120	121	122	123	124	125	126	14				
14						127	128	129	130	131	132	133	134	135	136	15				
15						137	138	139	140	141	142	143	144	145	146	16				
16						147	148	149	150	151	152	153	154	155	156	17				
17						157	158	159	160	161	162	163	164	165	166	18				
18						167	168	169	170	171	172	173	174	175	176	19				
19						177	178	179	180	181	182	183	184	185	186	20				
20						187	188	189	190	191	192	193	194	195	196	21				
21						197	198	199	200	201	202	203	204	205	206	22				
22						207	208	209	210	211	212	213	214	215	216	23				
23						217	218	219	220	221	222	223	224	225	226	24				
24						227	228	229	230	231	232	233	234	235	236	25				
25						237	238	239	240	241	242	243	244	245		26				
26						246	247	248	249	250	251	252	253			27				
27						254	255	256	257	258	259	260				28				
28						261	262	263	264	265						29				
						G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U

FIGURE 3.: This figure shows the numbering and indexing system for the 265 subregions, or "sites", in the study area.

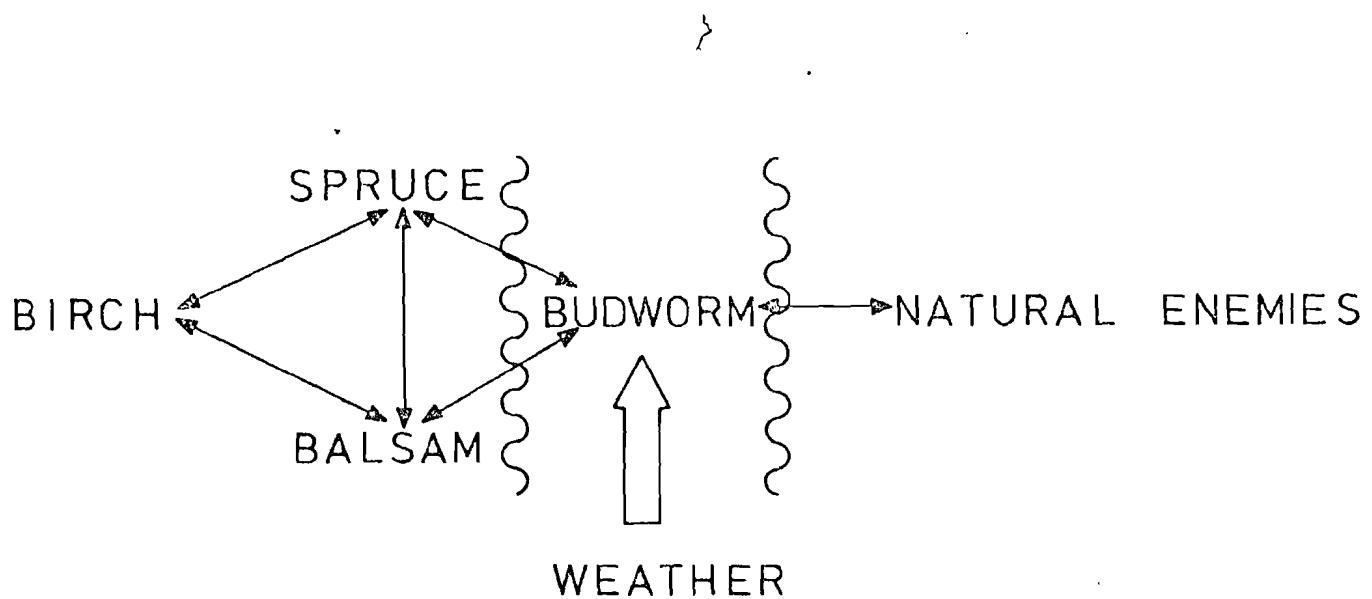


FIGURE 4.: The key roles or variables and their interrelations in the natural ecosystem. The principal tree species (birch, spruce and balsam fir) have a dynamic interaction of their own. This interaction is altered by the presence of budworm which consumes some spruce but primarily balsam. The budworm is in turn affected by a complex of natural enemies and a stochastic weather variable.

advantage between outbreaks and spruce and birch during outbreaks. The result is a diverse species mix.

As noted earlier, between outbreaks the budworm is rare but not extinct. Its numbers are then controlled by natural enemies such as insectivorous birds and parasites. But a key feature of this control is that there exists an upper threshold of budworm numbers which, once exceeded, allows the budworm to "escape" predation and multiply unchecked. There is, in other words a distinct but limited stability region at low budworm densities.

In addition to tree species and natural enemies there is a key stochastic driving variable, weather, which affects survival of the budworm and can flip the system out of the low density stability region if forest conditions are appropriate. Outbreaks cannot occur unless the forest has sufficiently recovered from the previous outbreak to provide adequate food. Even with the food conditions met, however, the budworm remains at low densities under control by natural enemies until the weather shifts to successive years with warm dry summers. Such conditions allow larvae to develop so rapidly that densities above the escape threshold are achieved. An outbreak is then inevitable, irrespective of weather.

In summary, the decisions on bounding the problem are as follows:

- | | |
|------------------------------------|--|
| Time Horizon | - 150-200 years |
| Time Resolution | - 1 year with seasonal causation |
| Spatial Area | - 17,000 square miles |
| Spatial Resolution | - 265, 6 x 9 mile sub-regions |
| Key Variables to Capture Behaviour | - ideally, three tree species, budworm, natural enemies and weather. |

This bounding of the problem immediately determines the number of state variables, which in turn affect the decisions about subsequent analytic steps such as optimization. Even though the previous steps of bounding seem to have led to a highly simplified representation, the number of state variables generated is still enormous. Table 1 summarizes, for this ideal condition, the minimum number of state variables necessary to represent the essential behaviour of the system in space and time. In any one subregion 107 state variables are required, but of course, for the whole 265 subregions a total of 107×265 or 28,355 state variables are required. Thus even this drastic simplification accomplished through the bounding exercise leaves an impossible number of state variables, thus demanding further simplification. It would, of course, be quite possible to develop a simulation model with this number of state variables. This would be expensive and time consuming to run and debug, but it would be possible. Our key goal, however, is to provide a useable and well tested model for exploring behaviour and policy alternates. With such a high dimensionality the model would become nearly as incomprehensible as the real world and the opportunities for systematic exploration would be greatly reduced.

As a consequence, a systematic series of further compressions and tests were made to determine whether the number of state variables could be significantly reduced. This led to four prime variables: tree density, foliage condition, budworm density and weather. The tree density actually had to be represented by 75 state variables associated with tree age but techniques were developed to collapse these into one for descriptive purposes. The effects of all the other variables can be incorporated implicitly so that this ultimate compression requires essentially four state variables per site or 1060 for the region.

The Model

The basic form of the model structure is shown in Figure 5. Budworm reproduction and survival, forest response, and control policies are independent for each of the 265 sites. Once each year dispersal occurs among the sites and the process is then repeated for the next simulated year. The budworm and forest response models were developed from the extensive set of data collected by Environment Canada over the past 30 years. Many of the component processes such as growth and reproduction have been examined in extensive detail by Morris et al. (1963) using multi-variate statistical procedures. But there are three critical processes that are not clearly understood at present and for which there is, at best, qualitative information. These three areas of semi-knowns are the effect of natural enemies at low densities of the insect, the detailed response of trees to defoliation and the specifics of dispersal. Since this problem of grappling effectively with substantive unknowns is central to any successful analysis of ecological problems, considerable effort was spent in developing a formal procedure to cope with such uncertainties. Moreover, these three particular areas of uncertainty are typical of many situations. Rarely, for example, is there much detailed information about events when the number of organisms is very small. Nor is there often much knowledge concerning very slow processes such as those involved in tree responses. Finally, dispersal occurs over such large areas that only the very recent application of radar technology has made it possible to define and quantify the form and magnitude of spatial contagion. And yet we know each of these three sources of uncertainty to be critical in determining aspects of renewable resource systems behaviour which have been particularly troublesome in past efforts of environmental management.

The ecological literature contains a particularly rich body of experimental and theoretical analyses of key ecological processes. These have led to the identification of classes of interactions, each

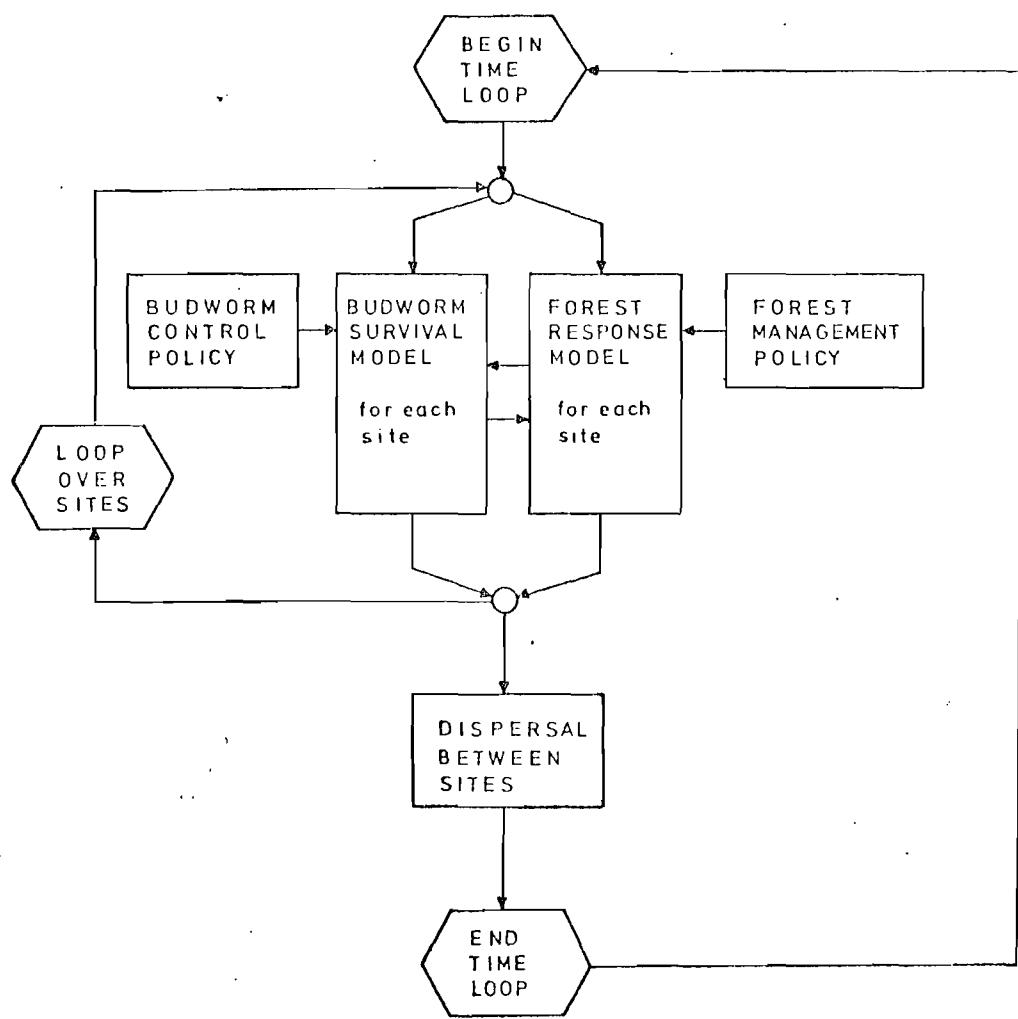


FIGURE 5.: The basic model structure for the budworm/forest simulation model.

characterized by a specific family of mathematical equations. For example, predators display a set of responses to prey or host species which fall into nine primary classes. Not only are each of these classes defined by a specific family of functional relations, but the biological attributes for each class have also been sufficiently well identified that quite qualitative information usually makes it possible to assign a specific example to the appropriate general class.

This work provides a theoretical framework allowing us to mobilize the existing information, however sparse, in a way which lets us proceed in a series of steps to gradually define a narrower and narrower range of possible relations. Having identified in this manner the classes of responses characterizing specific situations, it remains necessary to parameterize them. Even in the worst of circumstances, information usually exists to permit rough specifications of the parameters, leading to the definition of a maximum possible range for each response class.

The final step is to cycle these possible relationships through the full simulation model in order to define a "feasible range" of forms which result in retaining the known qualitative behaviour of the system. At that point informed judgement can informally select a "standard" relationship for use. Alternatively, an organized application of decision theory can assign subjective probabilities and so generate a range of possible outcomes.

The key point of this exercise is to directly face the reality of unknowns and to recognize that an organized approach in dealing with them can not only provide reasonable solutions that will allow the policy design process to proceed, but at the same time can provide very clear and specific priorities for future research.

A dynamic descriptive model of the sort described here is useless for prescription unless it presents opportunities for meaningful

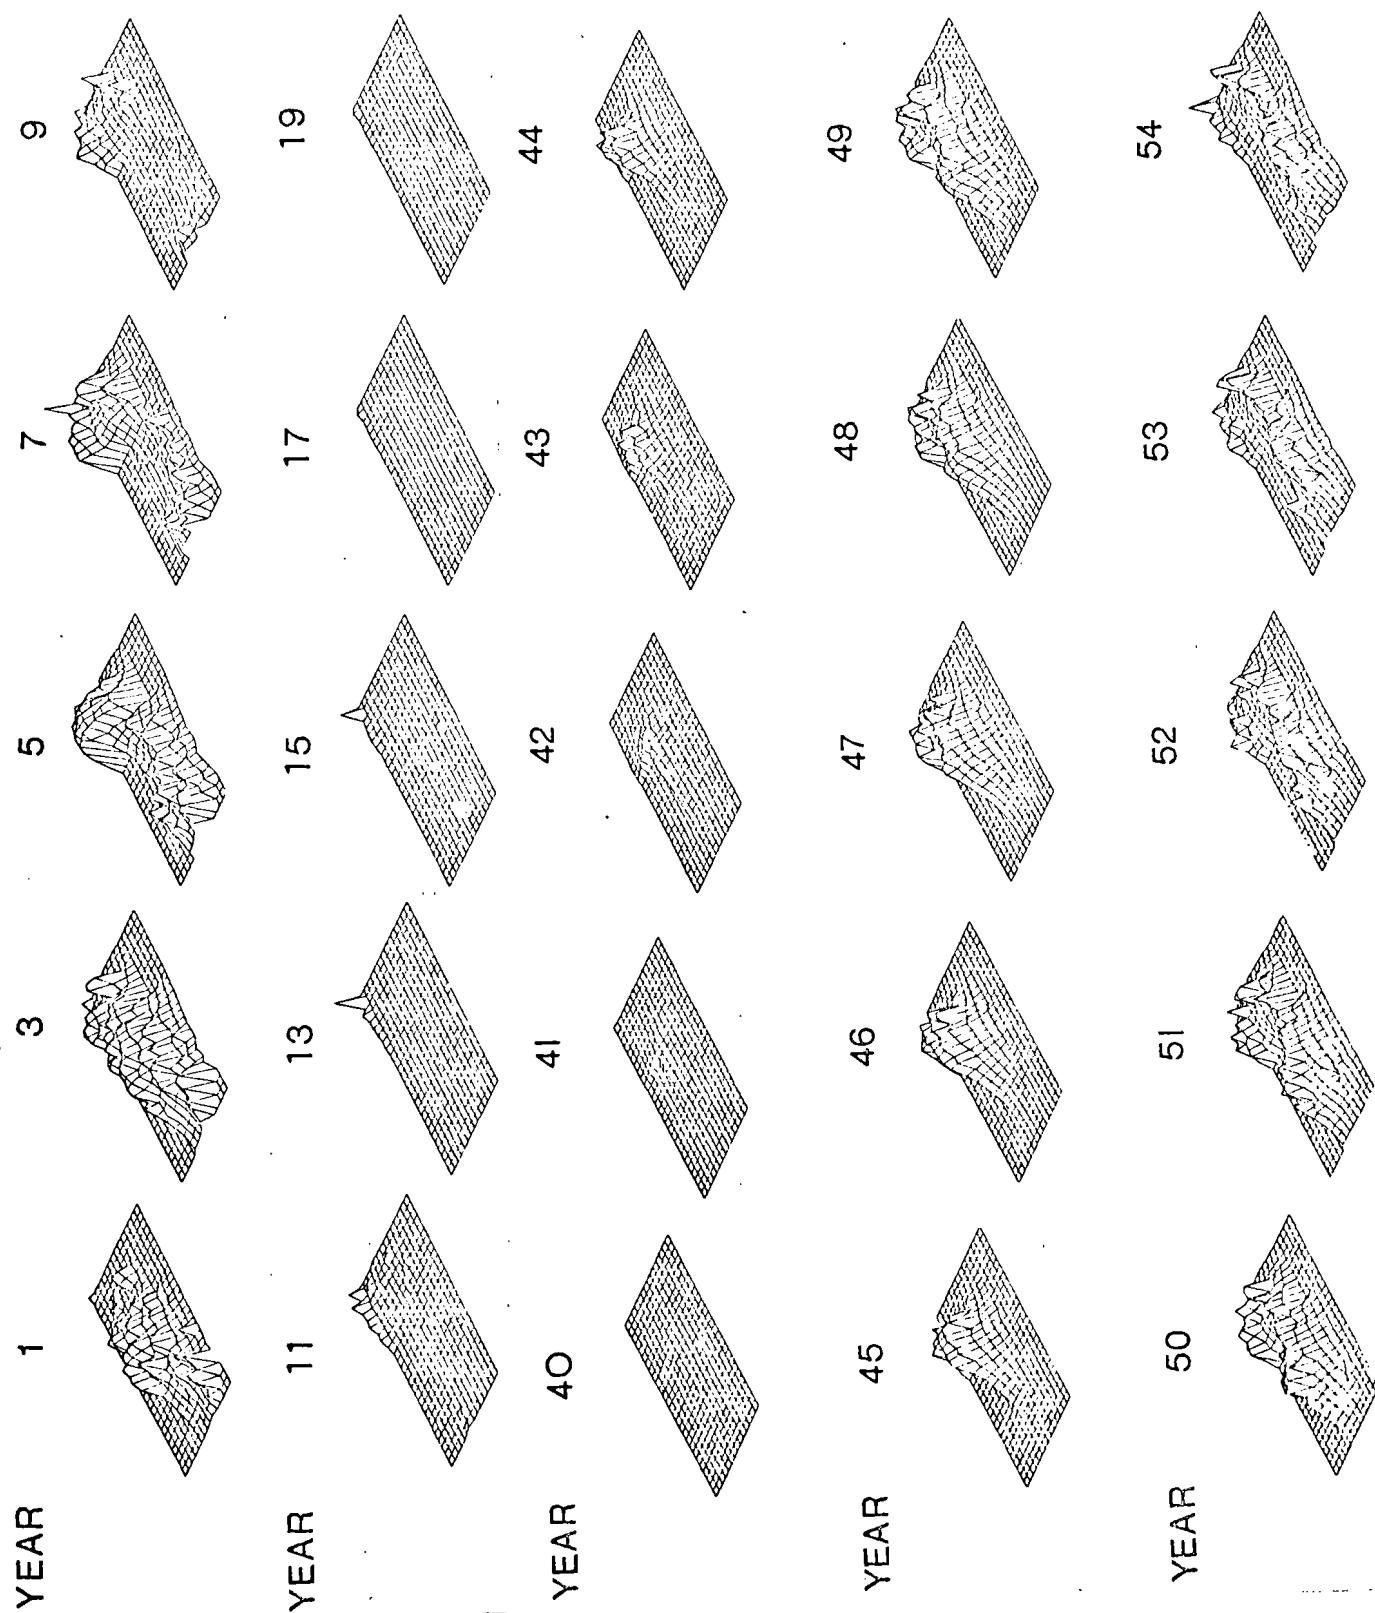
management intervention by policy actions. There are two main classes of policy action possible - one relating to control of the budworm and the other to management of the forest. These are structured in broad terms allowing for the exploration of not only insecticide control of budworm but for biological and other methods of control as well. Similarly the forest management policy can include specific actions of cutting by age in different regions of the forest and also, at least implicitly, a variety of silvicultural and tree breeding actions. Although the model is structured to accommodate a wide range of possible actions, for the purpose of this case study attention was largely directed towards budworm control using insecticides or bacterial agents and forest management using different techniques of scheduling cutting in space, time and by tree age.

Model Validation

Validation of an ecological model is always a difficult problem. In this example a statistically rigorous validation would require detailed historical information on all the state variables over a large spatial area and covering a very long period of time - at least 70 to 150 years. Only in that way could the full dynamic interplay of the system over time and space be adequately tested. The budworm case is rare in that such data are in fact available for a 30 year period, but that is scarcely long enough to instill a profound confidence in the model. Our present validation approach has therefore been to combine a quantitative comparison of state variable values over the period for which detailed historical data is available, with a qualitative comparison of gross behavioural properties (mean outbreak densities and variances, length of inter- and intra- outbreak periods, and so on) for the longer term. Some sense of this latter exercise can be drawn from Figure 6A. This Figure shows a series of computer drawn maps generated by the simulation model showing on the vertical axis the density of budworm

FIGURE 6A.: Computer Simulation Maps of Budworm Egg Density -
(No management).

-17-



eggs over a 54 year period. It covers two outbreak sequences and the essential pattern of these predictions have been confirmed by historical data from New Brunswick and elsewhere in eastern Canada.

In summary, the descriptive analysis led to the development of a dynamic simulation model that could describe the behaviour of the forest/pest ecosystem in space and time, with opportunities for intervention with a variety of management acts. It provides, therefore, a laboratory world in which the consequences of a variety of alternate policies can be explored and constitutes the essential base for the prescriptive analysis.

The Prescriptive Analysis

The goal of the prescriptive analysis is to provide a management tool which can aid in policy design and evaluation. There are three parts to the analysis as we implemented it. The first was the definition of a strategic range of management objectives, the second the application of optimization techniques to develop policy rules for each objective, and the third the development of a framework to broadly evaluate the consequences of each policy in terms of a wide range of potential management goals.

Strategic Range of Objectives

The uncertainties and unknowns in describing an ecological system are trivial compared to the ambiguities in defining societal objectives. The objectives that seem so clear at any moment can shift dramatically, as testified to by the recent concern for environmental issues. Moreover, as has been discovered by the water resource planners

in particular, even the best of policy analyses can founder on unrecognized or hidden public objectives. Since social objectives are hidden, ambiguous, conflicting, and otherwise indefinite, the analyses rarely can accommodate them satisfactorily. Hence they become uncomfortable, intrusive and divisive issues of confrontation. In response to this essential ambiguity of objectives, we felt it essential to identify a strategic range of objectives containing a systematic spectrum of plausible and not-so-plausible management goals. Any specific example drawn from that spectrum was initially considered only as a touch-stone and in no sense a realistic or desired objective. Again, our aim was to provide a tool for articulating and exploring alternatives, not a predictor of prepackaged social goals. The strategic range, as we conceive it, covers at one extreme objectives that attempt to achieve long term profit maximization and minimize probabilities of failure. At the other extreme, and equally unrealistic, the objectives seek more to retain the dynamic variability of the system in ways less sensitive to ecological surprise or changes in social and economic goals. These latter objectives attempt to achieve systems that are resilient, or robust; that work with the dynamic rhythm of the system rather than against it. If the first extreme represents the goal of a fail-safe world, the latter represents one which is safe in failure.

Five strategic touch-stones of this kind were defined as follows:

1. Unconstrained profit maximization.
2. Constrained profit maximization, where the maximum processing capacity of the existing logging industry sets constraints.
3. Recreation maximization, acting as an additional constraint on 2 above.
4. Budworm minimization, replacing the spraying policy of 3 above with alternate methods of forest management and budworm control.

5. Variability transformation, operating independently of 2 above, in which the goal is to transform the high temporal variability, which causes a boom and a bust situation for employment and the forest industry, to spatial variability. The goal in this extreme is to develop a forest ecosystem in which the budworm can be used as a forest manager and the essential dynamic interplay of natural forces is retained.

In addition, of course, two additional policies are explored - one of no management and one of the historical management. This produces a total strategic array of seven alternate objectives that, after evaluation and comparison, can be modified, combined and refined as starting points for a policy design dialogue with managers and specific interest groups.

Optimization

Given a range of objectives the next step is clearly to specify ways of combining available management actions in sets of policy rules appropriate for their realization. Simulation gaming, at an early stage, is a useful exercise for heuristically exploring the possible consequences of different management acts. Even after more formal optimization procedures have suggested specific rules, such gaming can still be a rich environment for dialogue. But the immense variety of different ways of combining acts in space and time demands more structured procedures as well. Figure 6B is an example of a gaming simulation run.

The approach we have taken is to regard the simulator as a means of bringing the real world into the laboratory. The various policies (whether obtained by common sense, or by common practice or

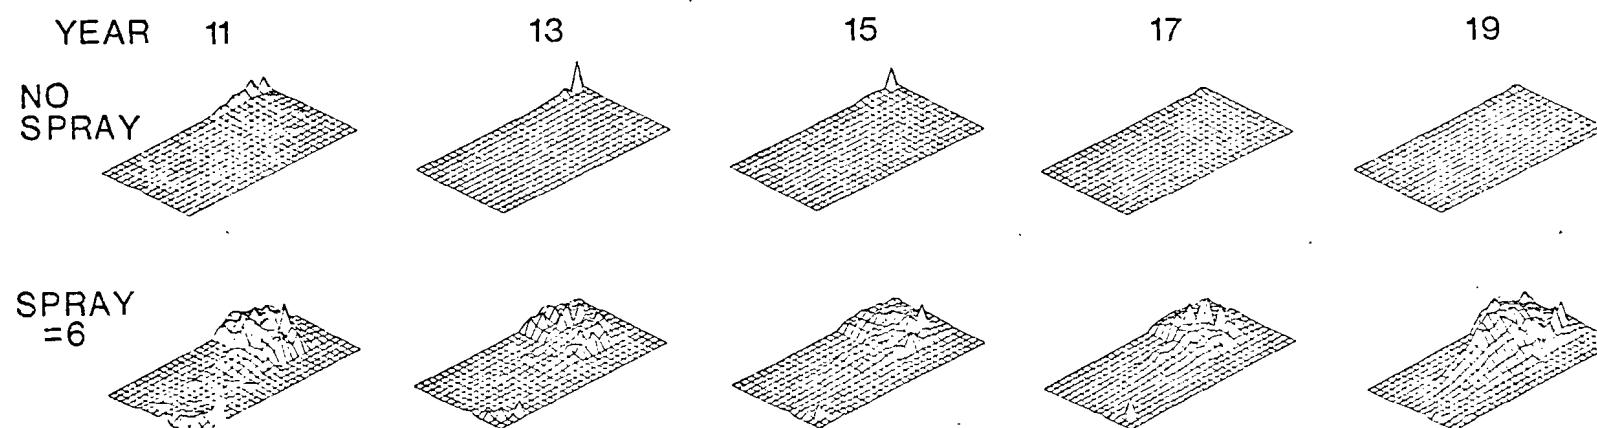
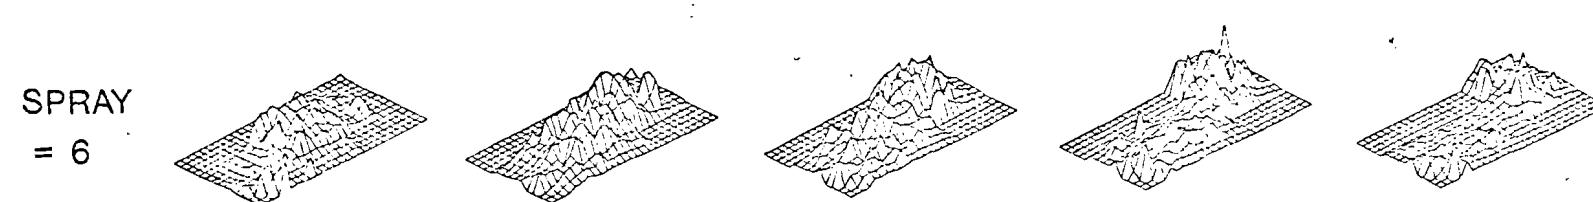
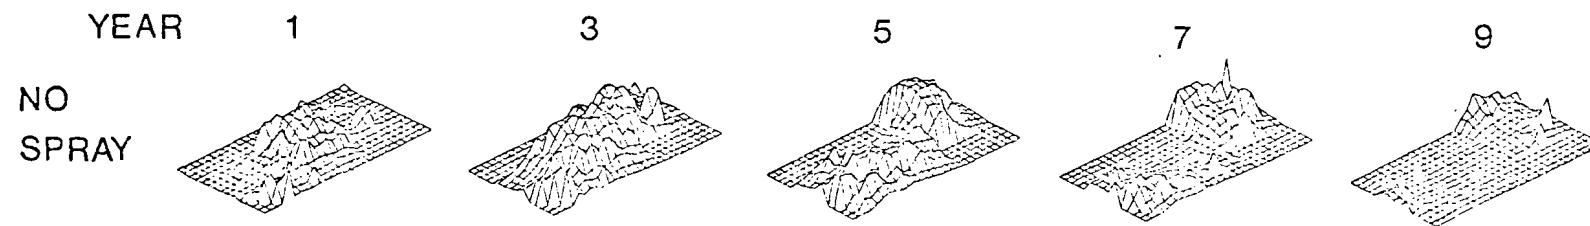


FIGURE 6B.: An example of gaming simulation in which a spraying procedure, similar to that used historically, produces sustained semi-outbreak situation.

through the use of an "optimizer") can always be compared by making a sufficient number of runs on the simulator. An analyst weak in analytic skills, poorly trained in the formulation of models, poorly informed about algorithms for solving classes of models, or unfamiliar with software availability may well opt to run many cases on the simulator to see if local improvements in a proposed policy is possible. Most simulation efforts unfortunately end up this way. Unfortunate because the high cost of using simulators to test many cases usually exhausts the patience and funds of sponsors to support development of an optimizer. If these funds had been used instead to develop a simplified model, then the process of determining an optimal policy for the simplified model could serve as a "brain" for the simulator and would have resulted in significantly better policies being found.

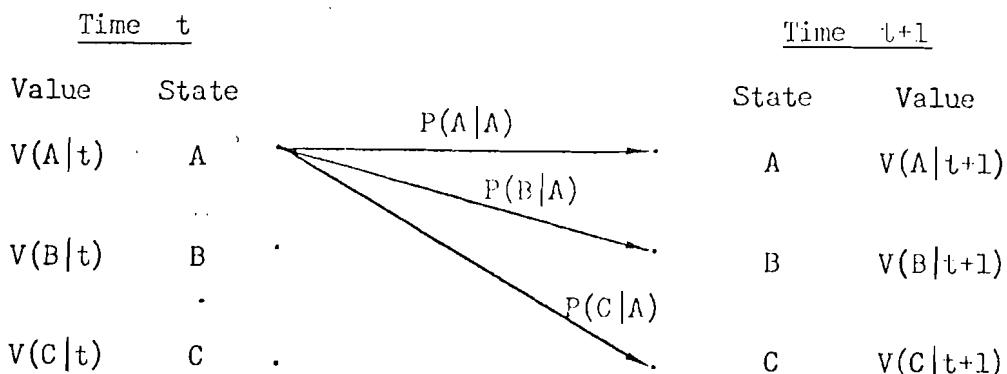
Generally speaking there are two types of analytic models that have had many successful applications: (1) "linear programming", and (2) "dynamic programming" models.

The first, the linear programming model, is characterized mathematically by a system of linear inequalities. Many kinds of non-linear relations can be practically approximated by such systems which can be both dynamic and stochastic. Software is available for solving such systems at reasonable costs even when they involve thousands of inequalities and variables.

The second, the dynamic programming model is characterized by a dynamic system that moves from any given state in time to the next without being effected by the past history of how it arrived at its given state. Many practical models can be cast in this form. In practice, however, applications are narrowly limited to those whose "state space" may be approximated by a low number of cases. In our research, however, we have pursued an alternative possibility - one which allows the state

space to be multidimensional and continuous in certain components. We were able to do this by finding a practical way to approximate the "pay off" for each state if one follows henceforth an optimal policy.

For the Budworm Optimizer we used a mathematical model closely related to the dynamic program -- the so called Markov Process. At each point in time t , the system is some state A, B, C, ... If in state A it will move to state A or B or C, ..., at time $t+1$ with probabilities $P(A|A)$, $P(B|A)$, $P(C|A)$, ... ; similarly if in state B it will move to A or B or C at time $t+1$ with probabilities $P(A|B)$, $P(B|B)$, $P(C|B)$, ..., etc.



In our application these probabilities can be changed at a price by engaging in certain alternative actions. The problem is to find the best choice of these alternative actions. This is easy to do if we know the value $V(A|t+1)$, $V(B|t+1)$, ... of being in various states at time $t+1$. Thus the expected value $V(A|t)$ is given by

$$V(A|t) = P(A|A) \{V(A|t+1) - C_{AA}\} + P(B|A) V(B|t+1) - C_{AB} + P(C|A) V(C|t+1) - C_{AC} \dots$$

where C_{AB} , for example, is the cost (revenue if negative) of transitioning

from A to B in time period t. If there are alternative actions in period t which can affect these probabilities, then the action that yields the maximum value of $V(A|t)$ is chosen. The procedure is thus a backward induction to time $t = 0$ but requires (in order to get it started) the knowledge of $V(A,t)$, $V(B,t)$, $V(C,t)$ for some future time $t = T$ in the future.

As noted a Markov type model was the one used for the budworm study. The key idea used to develop this analytic model was to view the single tree as an entity which changes state from year to year -- its state being defined by its age, stress, and the number of budworms it hosts. The tree, depending on the weather and whether or not it is sprayed or cut will (with certain probabilities) become one year older with certain stresses and egg densities or reverts to age zero and is replanted. If it were not for the spread from one timber stand to another of budworm eggs by the adult moth, this model has the merit that all other relations can be used with little or no simplification or change. This leaves open the question of how to approximate the effect of egg contamination. Several approaches have been posed and are dealt with elsewhere (Dantzig and Holling 1975). But acute simplification is the rule in dynamic programming applications, and we shall show how we have presently elected to live with this one after outlining the existing model and its solution.

For the simplified model we wish to find for every state (tree age, stress, and egg density) the optimal policy.

One way to determine optimal policy is to begin with a guess v_0 as to the entire discounted future value of a tree starting at age zero including the value of all its future harvesting and replanting (to time infinity) when we always carry out an optimal policy in the future with regard to the tree and its replantings. A tree planted a

year from now, has present value of $.95 V_0$ for its time stream from 1 year to infinity where 5% (say) is the discounted factor (without inflation). If for the moment we accept our guess V_0 , we are in a position to evaluate the present value of all other states. One begins by noting that as far as harvesting the lumber of the tree now (or in the future) it does not pay to allow a tree to become older than 60 years (say). If so then the optimal policy is to cut it down and its present value $V_{60} = V_0 + L_{60}$ where L_{60} is the value of the 60 year old tree as lumber (less any cost for replanting it). To obtain the V_{59} of a 59 year tree (which is in some state of stress and egg infestation) and, at the same time, to obtain the optimal policies: (1) cutting it down, $V_0 + L_{59}$; (2) leaving it alone, $.95\{pV_{60} + (1-p)V_0\}$ where p is the probability of the tree living; and (3) spraying, $-S + .95\{\bar{p}V_{60} + (1-\bar{p})V_0\}$ where S is the cost of spraying and \bar{p} is the probability of the tree living after it is sprayed. The policy which yields the highest value is selected as optimal. Note that the effect of random weather factors are part of the calculations (i.e., weather affects the probabilities of dying or the probabilities of moving from one state to another) so that values (and optimal policies) of various states can be determined backwards from the highest age 60 down to age 0. If it turns out that our guess V_0 checks with the value V_0 obtained by the backward calculations we accept it -- if not then we revise our guess up or down until it does check.

This procedure defines therefore an optimal way to apply the variety of management acts for a specific objective in terms of the values of the key state variables. These policy rules may be represented in the form of policy tables such as those shown in Figure 7. For any age of tree, foliage condition and density of insects the manager can either do nothing, spray (and the spray can be at different intensities and concentrations) or harvest. The advantage of such policy tables is that they are clear, unambiguous and can be easily applied by a forest manager attempting to manage a stand in isolation from the rest of the regional forest system.

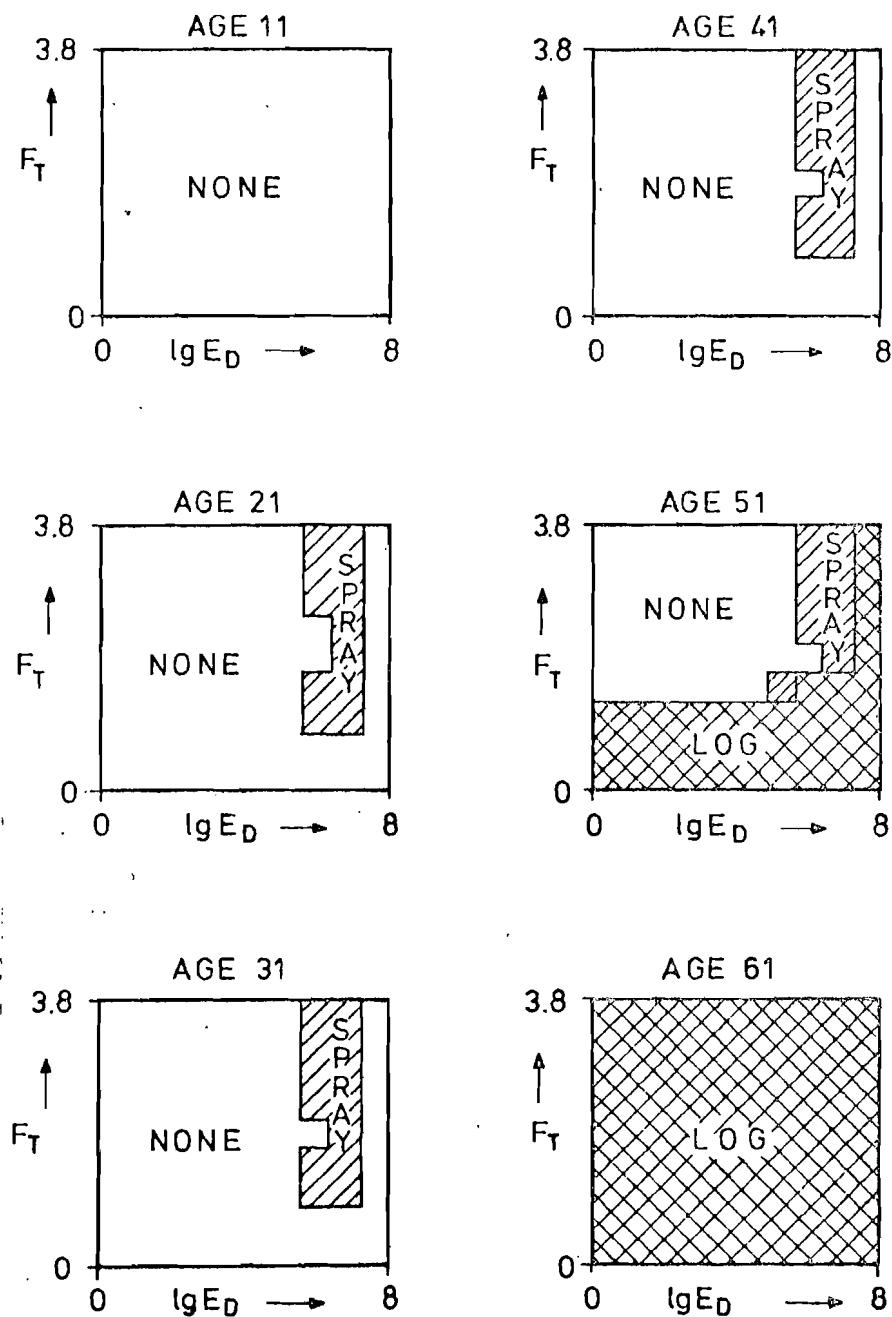


FIGURE 7: Policy tables for representative ages
(price = 55\$/cunit; $\rho = .05$)

But to achieve those "optimal" rules required gross simplifications due to the limitations of available optimization techniques. Two major simplifying assumptions were required. The first concerned a simplified expression of the objective function, and the second required that dispersal between spatial areas was unimportant. It was only in this way that the high dimensionality of the problem could be simplified to the point where dynamic programming could be successfully applied. Similarly gross simplifications will be required in most problems involving dynamic management of resource and environmental systems.

Dynamic programming is a particularly powerful and valuable tool for use in ecosystem management studies. But unless really substantial advances are made in its ability to handle certain classes of high dimensionality, it will properly remain a special-use "sub-optimizer" methodology only. For the foreseeable future, we will have to learn to make the most - without making too much! - of that.

Sub-optimal or partial optimal solutions have a useful role to play, however. The key to their constructive utilization is an ability to cycle such simplified policies through the full simulation model with all its complexity. By using a variety of indicators, each of these policies can be assessed in terms of a possible drift of solution from some broader societal and environmental goal. When this is detected then ad hoc, heuristic modifications of the policies can be employed to produce more desirable behaviour.

This process, again, should be in the form of dialogue with both managers and interest groups. As we said earlier, the optimization model was designed to provide a "brain" for the simulator. But that "brain" is a childish thing and for its proper functioning it requires the guidance that can only be provided by those that make policy and those that endure it.

Evaluation

A program of policy exploration and evaluation using a simulator requires the development of a rich array of social, economic and environmental indicators, and a framework for their use and interpretation. The manager must be able to converse with the model in a critical and flexible manner if the latter is to have any legitimate use as a policy design tool. A system of indicators was therefore designed to serve as the common, comprehensible language in which that conversation could take place. The grammar and syntax rules which give structure to the managers dialogue with the model are derived from the discipline of operations research and, more particularly, decision analysis.

The central difficulty in applying traditional decision analysis approaches to the budworm-forest management problem has been the essentially dynamic nature of both the system itself and the majority of possible policies for its control. The "present state" description of the system tells us only a little about future states, and the essence of a good management policy is precisely the ability to adapt quickly and successfully to the inevitable future state surprises, as they arise. We are really more interested in what the forest is doing than what it is at a given moment, and the standard paraphenalia of a discipline still mightily concerned with relatively complacent marbled urns has predictably been unable to give us quite the help we need.

Our ad hoc solution to the problem of specifying objectives for a time changing system has been to present the decision maker with full time series descriptions of the forest's behaviour, without regard to the policies employed to generate that behaviour. In principal, the choice problem simply becomes one of ranking time streams rather than state descriptions of the managed system. But the full panoply of time

stream indicators associated with any simple management policy is exceedingly complex. To enable consistent rankings of alternatives, we must first simplify, and simplify both drastically and meaningfully.

The first step in the process is straightforward. Each individual manager is asked to review the list of indicators, strike those of no or minor relevance to the determination of his time stream preferences, retain the rest, and add or alter where necessary. A representative list of one decision maker's "things I am interested in" is shown in Table 2.

A little bit of additional simplification can be done rather easily. These indicators are listed by intuitive groups of like "kind". It turns out that the decision maker's expressed tradeoffs among indicators within such groups are often independent of the values taken by indicators outside the group (a sort of preferential independence). As one put it "I can add apples and apples without caring too much about oranges". In addition, within-group tradeoffs were usually expressed as "noncompensatory" or threshold phenomena in which a given indicator became important only when it took a value outside a wide "normal" range. These convenient simplifications made it possible to reduce most of the Table 2 groups to single, aggregate indicators in a relatively unambiguous and intuitively plausible manner. An example is shown in Figure 8. Following similar procedures one manager reduced his initial long list of indicators to only three; one for economic effects (including logging, spraying, and operating costs), one for recreational value (including accessibility factors, forest composition, logging and insect damage, and so on), and one for social issues (essentially the level of labour force displacement due to forest destruction by budworm).

TABLE 2

Preliminary "Grouped" List of Relevant Indicators

A) Economic

- x_1 = "Profit" to logging industry
 x_2 = Cost of insecticide spraying
 x_3 = Cost of other control measures

B) Forest Appearance

- x_4 = Age class diversity
 x_5 = Proportion of trees in "mature" classes
 x_6 = Proportion of trees severely defoliated
 x_7 = Proportion of trees dead
 x_8 = Proportion of area logged

C) Social

- x_9 = "unemployment"; 1-prop. of mill capacity not filled

D) Forest Potential

- x_{10} = Amount of merchantable wood present
 x_{11} = Amount of merchantable wood harvested

E) Ecological

- x_{12} = Average concentration of insecticide in sprayed areas

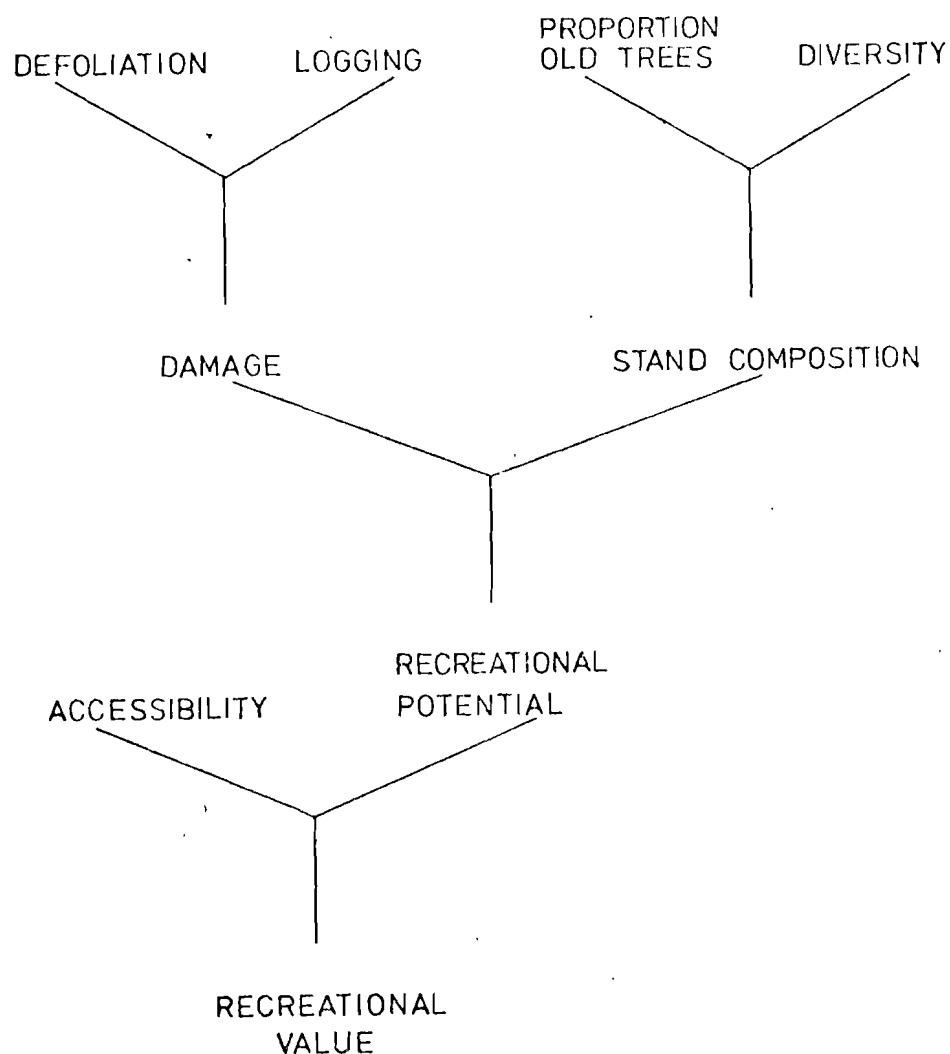


FIGURE 8: Steps in aggregating four basic variables into a recreational indicator.

Having completed the initial indicator selection and aggregation, we possess a reasonably concise way of describing any given pattern of system behaviour. The task remains, however, of systematically, meaningfully, and unambiguously ranking alternative sets of time streams such as those shown in Figure 9. And most regrettably, this seems to present a problem wholly beyond the capacity of present theory and methodology in decision analysis. Consider for a moment, the difficulties.

Time lies at the heart of all our problems here. If we wish to assign a single ranking value to a given set of indicator time streams, we must ultimately compress indicator values across time. The first inclination is to take variously weighted time averages of the indicators; means, discounted sums, and so forth. But any such time averaging scheme implies an explicit attitude of intertemporal tradeoffs through which we are willing to relate the future to the present. It is arguable whether standard "1-15%" discounting arguments can be defensibly applied to even purely financial matters. Their appropriateness outside the world of capital investment is highly suspect, to say the least. And despite the large volume of writing on "social rates of discount", little of practical import has yet been said on this matter either. It would seem that our society has yet to agree upon a fixed rate at which it is willing to discount its posterity into insignificance. And whatever it may be, the proper solution to the discounting problem certainly does not involve convincing them to do so.

Even if the overall time averaging problem could be resolved, however, we are left with the problem of assigning appropriate ranking weights to different temporal patterns of an indicator. Surely an unemployment rate time stream averaging 10% and based on alternating years of full employment and 20% unemployment deserves a different ranking from

one which holds a static 10% year after year. Potentially meaningful properties of time stream behaviour are almost certainly captured in correlation and runs statistics as well as variance estimates. As with the averaging problem, the issue is not whether we can perform requisite calculations - which are trivial - but how we could make the results meaningful to the manager - which is not.

One potentially useful compromise approach to the time problem has been to compress indicators across kind but not time, resulting in a single aggregated value function time trace (Figure 9). The key here is conditionally to assume temporal independence of indicator tradeoff values, in essence pretending that the relative weighting attached to various indicators in a given year is independent of their values in neighboring years. With this assumption made, intra-temporal, inter-indicator tradeoffs are evaluated using a standard multi-attribute, revealed preference approach and an overall value function calculated. This function is applied independently to each year's separate indicator values, generating the aggregate value time stream shown in Figure 9. In our work we have permitted no discounting of the component indicators or final value stream because of the ambiguities inherent in aggregating across differentially discounted indicators.

The manager has at this point reduced his ranking problem to one of comparing a single aggregate value stream for each pattern of system behaviour in question. By visually pairing these value streams with their component, indicator streams, the decision maker may be able to interpret the former consistently and to evolve a stable ranking pattern. The pairing also serves to show, through the component indicators, which portions of the value stream are most likely to be sensitive to the temporal independence assumption. We have no methodology to cope with this sensitivity but the "flag" at least serves to temper our interpretation of the aggregate value stream with skepticism.

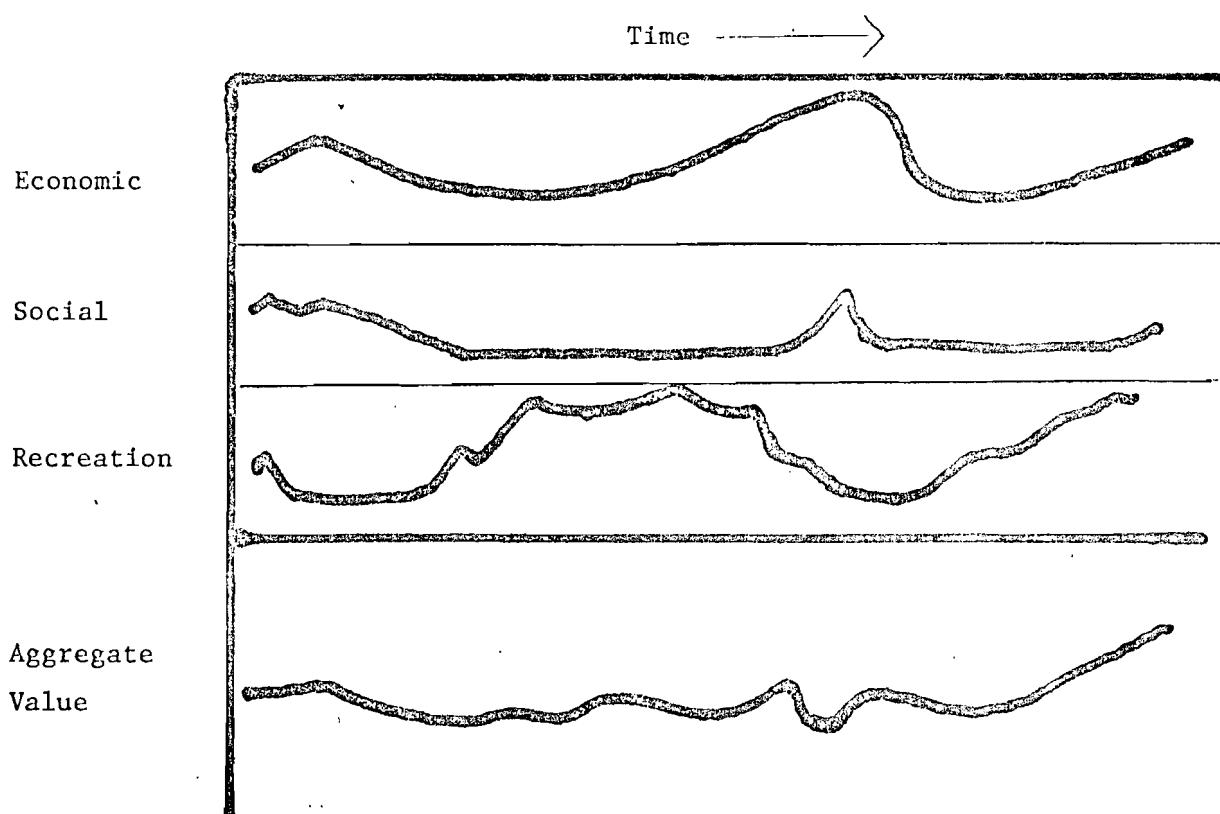


FIGURE 9.: An example of time traces of three selected indicators and their aggregate.

Further work is underway at IIASA this summer which seeks to improve upon the present unsatisfactory state of affairs. It appears unlikely that any breakthroughs will occur in the areas of discounting or, generally, handling intertemporal tradeoffs, but some progress on crucial issues of communication may be expected.

SUMMARY AND CONCLUSIONS

What Has Been Done

The intent of this case study was to see just how far one could proceed in combining the best of ecology, modelling, policy design, policy evaluation and decision theory, towards a realistic and characteristic problem of ecosystem management. The key ingredient of this analysis was the development of a rigorous, parsimonious and well-validated simulation model that explicitly addressed issues of unknowns and uncertainties. It provided the laboratory world for the development and exploration of alternate policies. The process led to optimization, where the limitations of existing techniques required an evaluation of "sub-optimal" policies through the simulation model. This process of evaluation in turn generated the need for a rich variety of social, economic and environmental indicators that could be used to judge the consequences of alternate policies.

Although it has not been mentioned, it additionally became essential to develop an alternate array of indicators, explicitly designed to handle the uncertainties or unknowns. These are necessary because ecosystems and, for that matter, social systems, are generally multi-equilibria ones in which each equilibrium is bounded by a stability region. Very little information can usually be mobilized to concretely

and specifically establish where those boundaries lie, or how the stability regions may contract with the application of management activities. There is, however, growing evidence from fisheries, forest and other ecological systems which suggests that these domains of stability do contract with no obvious indications until collapse occurs. Hence three classes of "resilience indicators" were also developed - one set measuring the unused environmental capital that would provide the alternate options required in the event of an unexpected event, one set relating to measurements of the stability boundaries, and one set concerning the resilience of benefits. The latter are generated by explicitly simulating specific kinds of policy failure and monitoring the stream of benefits thereafter.

The final key piece that linked the whole range of techniques was an explicit effort to generate a strategic range of objectives as first-cut management alternates. This range was designed to cover both non-resilient and resilient objectives with the intent of providing a rich menu for comparison and subsequent modification.

But so far we have emphasized mainly the techniques used, saying very little of the process by which they were developed and employed in the research program. And yet a crucial element of this exercise has been the inter-disciplinary and inter-institutional character of the operation. From the start the Maritimes Forest Research Centre, Department of the Environment, Fredericton, was involved in every stage of the project. Much of the economic and ecological analysis was performed at the MFRC Laboratory in close conjunction with other members of the team. The Institute of Resource Ecology, University of British Columbia, provided the expertise in systems ecology and mathematics. The whole activity was given focus and a disciplinary breadth at the International Institute of Applied Systems Analysis, Vienna. The final group of collaborators involved ecologists, systems mathematicians, and operations

researchers, covering a wide spectrum of talents. Despite that interdisciplinary breadth and the several thousand miles between the three key participating institutions, the degree of cooperation and communication was truly remarkable, in large part due to the initial fostering and flexible interactive environment of the International Institute of Applied Systems Analysis.

What Is Missing

If the test of this study is whether critical systems analysis can provide a more effective approach to ecosystem management problems than past approaches the answer is yes. At the very least, a quite specific list of research priorities can be defined, critically focussed on management needs, and leading to a more effective expenditure of available funds. Similarly the exploratory policies generated, although needing further modification as the effort leads to implementation, suggest management routes to greater benefits and robustness at considerably less cost.

But if the test is whether this range of techniques and the new ones added are adequate to the problems at hand, then the answer is no. Several major issues are quite unresolved. One of the most important concerns the difficulties of meaningfully aggregating indicators across kind, time, and space so that rational preferences can be expressed among alternate futures. The use of present discounting procedures to handle intertemporal tradeoffs is clearly totally inadequate, the more so because it makes the problem deceptively tractable. We know of no meaningful and effective way to make this time compression. Similarly fundamental problems in the effective application of optimization techniques have already been discussed.

Beyond the technical inadequacies, however, there are more important elements totally missing from the full process of policy design. When we compared the steps taken in the present case study with others that were developing simultaneously at IIASA - a total set of analytic steps began to emerge which would cover the full range of activities required in effective policy design. These are shown in Table 3 and in Figure 10. The budworm case study concentrated on the steps contained in the heavily outlined region of the table. Our implementation phase is just starting and it will encounter major issues of practical concern relating to availability of an infrastructure of roads, logging logistics and capital availability. That effort will need to develop additional tactical and more detailed simulation models and optimization routines and must proceed in close interaction with those agencies and industries actually responsible for management. The step of implementation is one which very few exercises of systems analysis have successfully accomplished, and certainly there are few major examples in the environmental field, outside of water resources work, which have involved effective implementation. That will be the challenge and acid test for the budworm programme over the next few years.

An equally important missing element in our programme has been the embedding of whatever policies are developed within the larger socio-economic reality of New Brunswick and of Canada. Typically, relatively little effort is expended in specifically addressing these questions of social embedding. And yet it is essential to develop some general overview of the broadest consequences of local policies. A promising lead has been provided by the energy case study at IIASA, which has argued that it should be possible to develop some simplified alternate societal models that in no sense pretend to be accurate representations of reality, but rather provide a framework or "mythology" to interpret the consequences as one person's view of the world. But in its final analysis, those societal consequences must be explored by those who must make the management decisions, and those who must endure them.

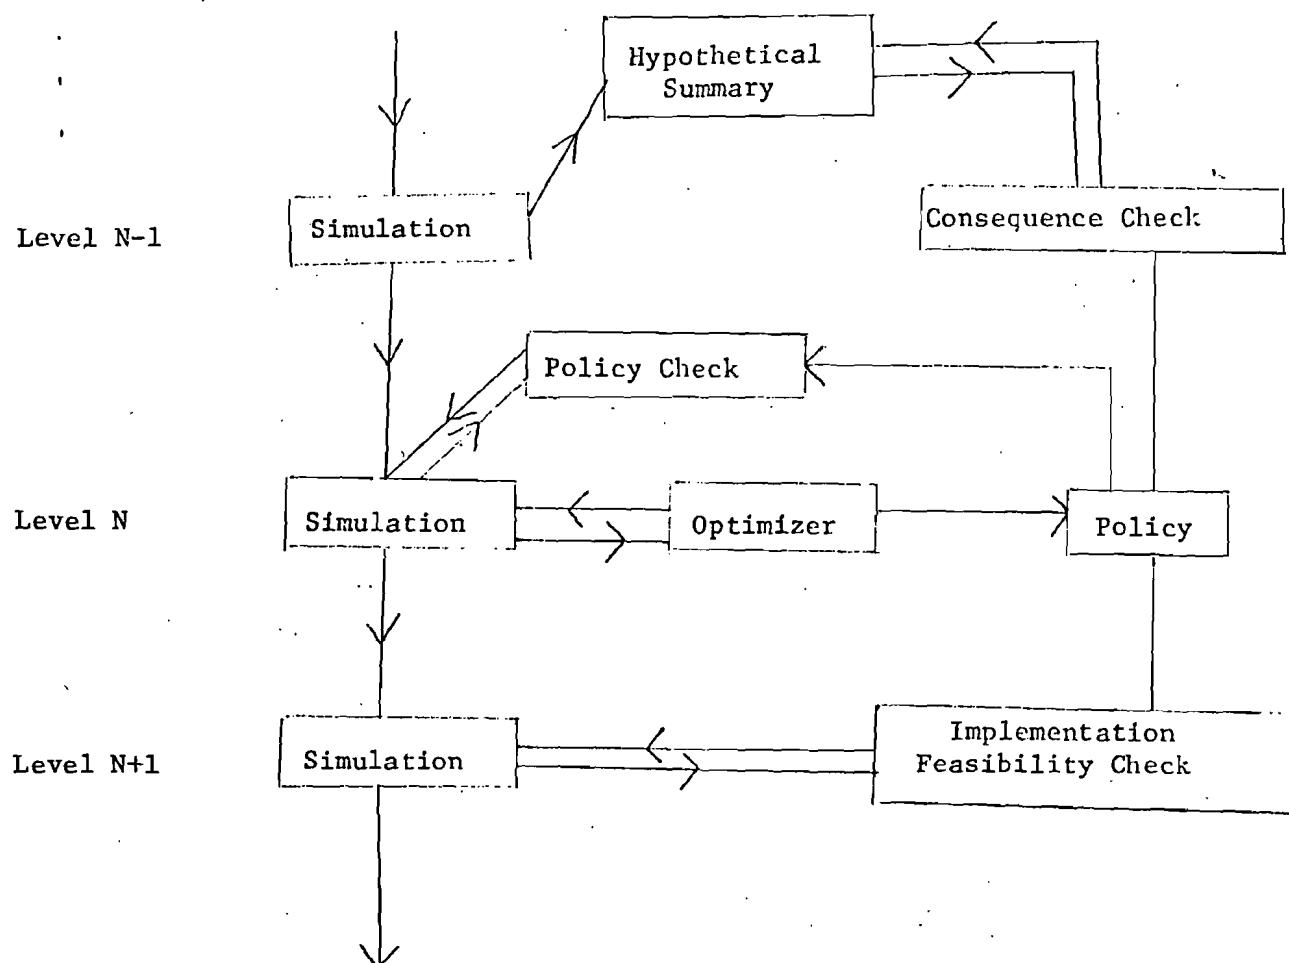
TABLE 3: THE FULL POLICY DESIGN PROCESS

Analytic Step	Systems Level	Function	Technique	Purpose
Hypothetical overview (embedding)	N-1	Consequence Check	-	To assess larger societal consequences of local policies
Detailed Dynamic Description	N	System Description	Simulation Model with full spatial disaggregation	To describe, dynamically the local system with enough confidence in its reality to be treated as reality.
Policy Prescription	N	Policy Design	Strategic range of objectives; optimization using simplified simulation model	To develop policy rules which are <u>only</u> state dependent, using simplifying assumptions.
Policy Evaluation	N	Policy Check	Cycling policies through full model; generating indicators (social, economic, recreational, environmental and resilience); decision Analysis	To evaluate broader consequences and feasibility of policy rules within the local system.
Policy Implementation	N+1	Implementation Feasibility Check	-	To develop detailed operational rules for implementing policy.

FIGURE 10: Flow Chart for Policy Determination

Level 1

Level 2



The effective exploration of societal consequences is dependent upon communication, and communication in its broadest sense. That too is an area of neglect in this and nearly all other efforts. If we were to devote 5% of the ingenuity we now spend on analysis to innovative and effective ways of communication, the payoff in terms of improved management would be astounding. The question at issue is how information concerning such a complex system can be presented in forms that are clear, understandable and usable; usable in a way that the perceptions and experience of the non-experts can be brought to bear on the analysis in a full and effective manner.

One final point needs re-emphasis. Past efforts in resource management have been essentially trial and error approaches to coping with the unknown. And indeed that is the way our society has advanced since the industrial revolution. Existing information is mobilized to suggest a trail, and if an error is detected then that provides additional information to modify subsequent trials. But we are now at the point where the intensity and extensiveness of our trials generate errors that are potentially larger than our society can afford. Trial and error seems increasingly to be a dangerous method for coping with the unknown. We need a new strategy to deal with ignorance. The concept of systems resilience provides at least a hint of a direction to proceed, focussing as it does not on the prediction of future surprises, but on designing systems that have the internal resilience to absorb those surprises when they inevitably appear.

Unless we can integrate into our design activities some such approach for dealing explicitly with the unknown, unless we honestly and effectively address the larger issues of social embedding and meaningful communication, all that we applied systems analysts can do is promise larger disasters, achieved faster and in a more pretentious and disciplined manner.

REFERENCES

Dantzig, George B. and C.S. Holling, 1975. Determining Optimal Policies for Ecosystems. In preparation.

Holling, C.S. et al., 1974. Project Status Report: Ecology & Environment Project - IIASA Report SR-74-2-EC.

Morris, R.F. et al. 1963. The Dynamics of Epidemic Spruce Budworm Populations. Mem. Ent. Soc. of Canada, 31.