A FORMALISM FOR MANAGEMENT OF SURPRISE
or
HOW I LEARNED TO DESIGN DAMS AND TO HATE SYSTEMS ANALYSIS

M. Fiering

June 1974 WP-74-20

A Formalism for Management of Surprise

or

How I learned to Design Dams and to Hate Systems Analysis

M.B Fiering

After many centuries of designing engineering structures and systems within a deterministic framework, it has become fashionable to deal explicitly with uncertainty as an important component of planning and design strategies. Advances in applied statistical decision theory, coupled with the wide availability of computing machinery, are at the root of this transformation, and the recent literature is replete with studies of systems, large and small, under various conditions of uncertainty. This paper deals with a few rules for decision-making under a special category of uncertainty--namely that associated with the occurrence of events which could not be foretold, let alone assigned a prior probability of realization within a given design horizon.

The use of liberal factors of safety has a long history in engineering design; it is commonplace to be derisive about these factors, and to call them "factors of ignorance" or other less endearing terms. But this is not entirely fair, because it has been traditional to have the safety factor reflect the degree of uncertainty inherent in the design and the cost (or danger) associated with failure. We speak here of the more obvious modes of failure; these include structural failure or collapse (of a building, a dam, an hydraulic control line) and operational failure (inadequate flood storage, inadequate irrigation supply, etc.). Thus we note that structural safety factors are of the order of 1.5 or 2.0, while the safety factor against failure of an earth dam by sudden drawdown, based on extremely conservative assumptions, is around 1.1. It is impossible to assign specific numerical factors of safety against hydrologic extrema, but we try to identify flood frequency characteristics and design against an event characterized by a specific return interval.

To these elemental considerations of uncertainty we must add a few new classes. Suppose we have at our disposal two decision variables, x and y, and that we seek those values (or that decision), say x^*, y^* for which the system response $f(x^*, y^*)$ is optimal. Typically the function f is some measure of net benefits or the benefit: cost ratio, appropriately discounted. The decision variables x,y are generally not free to range over all possible values but they, or some functions of them, are constrained by the conditions of the problem. Thus the derivatives of f with respect to the decision variables are not necessarily zero at the optimum.

We consider Figure 1, the contour map defined by the loci of equal response, or functional value f, on which it is desired to locate the decision (or x,y couple) where the response f(x,y) is maximal. Under the first class of uncertainty we decide on the values x,y(to be built) and undertake construction. In this section, we use the terms build and construction as if all decisions were structural components. But this is merely to avoid the ungainly terminology associated with repeating each time that a decision can be an operating rule or management decision, not merely a structural measure or capacity. For a variety of reasons relating to structural inhomogeneity, unreliable quality control, communication or human errors, etc., the finished system is characterized by a design different from the scheduled couple x,y; we call this $(x+\Delta x, y+\Delta y)$, as shown on Figure 1. This is tantamount to a small movement in decision space, but the contours of system response remain unchanged. This class of uncertainty is traditionally treated by application of a factor of safety.

The second form of uncertainty is the target of much of the massive effort in stochastic modelling of systems, particularly those which purport to represent environmental, ecological, meteorological and socio-economic interactions. It accommodates the fact that system components (typically classified as inputs, controls, demands and outputs) are rarely known deterministically. For example, streamflows, population projections, economic demand functions, the discount rate(s) and ecological processes can be estimated more or less well, and the importance of devoting explicit attention to their variations or instabilities dictates the extent to which stochastic considerations must be built into decision models. In the notation of Figure 1, we specify the decision x,y but realize a response

$$f'(x,y) = f(x,y) + \Delta f(x,y)$$

where the increment $\Delta f(x,y)$ measures the departure from the surface $f(x,y)$. The magnitude of this departure depends on random influences effective at the particular couple or decision vector x,y. In other words, the surface $f(x,y)$ is replaced by a mantle of variable thickness, with those responses highly susceptible to random fluctuation associated with thick mantles within which the actual system response might reasonably fall, while more predictable responses lie within closely contained mantles. We conceptualize the addition of at least one additional dimension to the system description; this dimension subtends some deterministic scalar response. But if response includes random fluctuations, then at least one additional dimension is required to describe the variation. However many dimensions are utilized, it is clear that the response surface itself remains fixed and that the realized outcomes migrate among the cloud of points which define the uncertainties inherent in system performance. As in the first class of uncertainty, we try to deal more precisely with where a particular realization or outcome will reside, contingent on a set of alternative responses and on some information concerning our ability to describe or even define the relevant random processes.

We now move to the third, and most interesting, class of uncertainty. It differs from the first two in that the response surface changes after the decision x,y is implemented. Many examples can be drawn from ecological experience; a classic case is that of the use of DDT. After years of widespread application, the "rules of the game" were abruptly modified and the response surface associated with the decision to spray was drastically changed, reflecting important damages and losses. Another case is the occurrence of a major environmental accident...a chemical or oil spill, a nuclear accident, a pollution episode of one sort or another...which causes the ecological system to "flip" (cf Holling

and Fiering and Holling: Management and Persistence of Perturbêd Ecosystems, <u>IIASA Ecology Project</u>, 1974) from one domain of stability to another. In other words, with reference to Figure 1, a whole set of new contours is dealt and the system is evaluated under a new regime, or under new criteria.

It might be argued that the consequences of such "surprises" could be reduced by the collection of more and better data and better understanding of the natural order of things. Indeed, a calculus has been advanced for specification of optimal data bases in special problems. The point here is that systems applications in important areas of human endeavor invariably deal with significant information gaps and uncertainties; moreover, no foreseeable models, no incipient insights, will reduce these uncertainties and gaps to levels which completely preclude surprise. The consequence of these gaps is that inconsistencies can enter the decision-making process; we propose here a strategy for dealing systematically with them.

At some earlier time in the history of technology the issue might not have been so serious, but our society is being inexorably driven toward problems of a larger scale, toward global considerations, toward scientific and technologic interventions and commitments which for all practical purposes are irreversible. We cannot hedge much longer, for example, with respect to generation of primary energy or its ultimate distribution through secondary and tertiary networks, even while legitimate environmental interests press for more rigorous pollution standards and better enforcement. The arguments are compelling, conjuring images of generations yet unborn, of denuded forests and of imbalances far more serious than the mere destruction of a particular piece of wilderness. It is not appropriate here to entertain the meta-physical arguments concerning the extent of our responsibility toward these future generations, or to interject judgments on whether or not we are so powerfully committed along a trajectory of consumption that preservation on our planet of life as we know it represents a feasible target. The inescapable facts are that we are galloping toward decisions which refuse to be delayed, that we will never have enough information to be perfectly comfortable about having to make them, that they have such long lead times for implementation as to be essentially irreversible, and that they are too expensive to initiate parallel tracks which allow for some maneuvering room. Part of the information basis for judging these decisions is the extent to which we, or our progeny, might be <u>surprised</u> by their consequences. We seek a <u>calculus of surprise</u> which can be utilized, with some of the more traditional metrics, for evaluation of program options. An example of surprise in a non-ecological setting is the recent history of U.S. oil policy and its consequences. For generations Arab disunity dictated reliance on the security of oil supplies to the U.S., and it seemed that contingency plans need not be made. But a measure of Arab unity was achieved, and however good the U.S. "system" for generating and distributing energy from oil, it performed badly under the changed rules. It was, and presumeably still is, inflexible, strongly subject to surprise; the long gasoline lines bore testimony to this.

These questions can be paraphrased in the ecological terms introduced by Holling. How resilient is the proposed system? How great is its capacity to absorb unanticipated perturbation and to continue usefully to function? How brittle is its optimum? Can it roll with the punches? Can it persist under environmental stresses whose magnitudes and frequencies cannot be foreseen? Can we trust our system to withstand stresses whose origins are now, and surely will remain, unfathomable?

This paper addresses a design formalism for systems which must operate under threat of extrema, including those events for which estimates of subjective probability can reasonably be made (e.g., extraordinary floods beyond the worst flood of record, the carcinogenic effects of cyclamates, etc.) and events which cannot be defined, let alone associated with some level of probability. For example, we could not reasonably have predicted *a priori* the now well-known effects of DDT, nor could we have agreed on a probability density for various intensities of these effects even if some perceptive biochemist had sounded the alarm. Moreover, no clear policy could have emerged simply by documenting the ecological threat; the trade-off between damage (particularly limited damage) and starvation (assuming for the moment that there is no immediately available alternative for pest control) is very elusive, and certainly depends on whether the decision-maker is starving, prepared to augment someone else's depleted crop, or merely looking on from afar.

Holling and others have remarked that our knowledge of ecosystems, however extensive, will always be exceeded by our ignorance. Thus we will always run the risk of being surprised by environmental consequences, and a traditional factor of safety, at least in the structural sense, is inadequate protection against this form of surprise. We thus distinguish between calculated risk, however that calculation might be made, and surprise. This difference is more profound than the familiar distinction between risk and uncertainty. We deal here with events which are not defined, not merely with those events for which we cannot reasonably assign probabilities. We plan to promote resilient systems, to discourage brittle ones. We plan to explore the region of the response surface near the optimum, to determine what happens if the system "flips" off its peak and tumbles into the surrounding lower region (How steep? How far down? How fast?) And we plan to investigate what happens if the rules are changed to the extent that a new deck of contours is dealt. We posit that the peak of the response surface may not be the best place to be because it may be so situated with respect to the boundaries of the domain of stability that a surprise will drive the system beyond its stable regime into a new, uncharted domain.

An Example

Consider the water-resource system in Figure 2. Two up-stream reservoirs service in-stream water demands in accordance with the standard or Z-shaped operating policy in Figure 3. The policy is characterized by two parameters, the reservoir capacity and the target draft, and by the assumption that the total inflow for the current time period (day, week, month, season, year or whatever) is known at the start of that period. This appears to be very restrictive, but experience over many years suggests that reservoir inflow and outflow are continuous variables and that the characteristic time period for most models can be made small, thereby rendering the assumption acceptable. The abscissa of Figure 3 gives the total amount of water available, consisting of initial reservoir contents plus inflow during the period. The policy ordains that if this is not greater than the target, the total supply is released and the reservoir remains empty. Any available supply in excess of the target is stored until the capacity of the reservoir is reached, whereupon the reservoir spills unavoidably.

Under the initial objectives promulgated for this simple system, each reservoir services its associated target without regard for the other reservoir or for the potential use to which the water might be put by the city located downstream, but there is a penalty function for not meeting a downstream target expressed at the city. Releases from the two reservoirs are assumed to be additive with respect to the downstream target, and there is no intermediate or unregulated inflow entering the system between either reservoir and the city. Thus if one upstream target were violated, the downstream target might still be met by a spill from the other reservoir. It is convenient to think of the upstream targets as in-stream uses, but this need not necessarily be the case. The reservoirs could be used to meet irrigation targets on the assumption that the return flow over a long time period, say a year, were equal to the diversion. The point is not to quibble over the exact uses of the water but to investigate system performance under this and a new set of system objectives.

We assign benefits to the several releases. The numerical values of flows, storages, targets and capacities are all integers to facilitate computation in this example, so that costs and benefits are then readily tabulated for each of the few possible draft and capacity combinations. The inflows are presumed to derive from a Markov process at each dam site, with the relevant data shown as part of Figure 2.

It is then a simple matter to calculate the steady-state inflow distributions at each site, the steady-state storage probabilities at each reservoir, and the draft probabilities at each site and at the city. From these and from the simple benefit functions we compute expected net benefits, the benefit:cost ratio (discounted and un-discounted) and a few miscellaneous summary statistics for system operation. There is nothing extraordinary about this exercise; it has been done, for one reservoir, by Thomas in 1958 (Harold A. Thomas, Jr: unpublished memorandum to the Harvard Water Program), by Fiering (for correlated Gaussian flows) in 1961 (Myron Fiering: Queueing Theory and Reservoir Design JASCE, Hyd Div), and by others since. It is a straight-forward matter to locate the optimal design (or combination of targets and capacities), albeit it is a tiresome computation. Suppose each reservoir can be as large as 3 volume units (4 choices) and that the annual flows cannot exeed 4 units. It is then sensible to talk of annual targets of 1 or 2 volume units, so that the total number of design combinations, at both reservoirs is 4 x 4 x 2 x 2 = 64. One of these is optimal, as shown in Table 1.

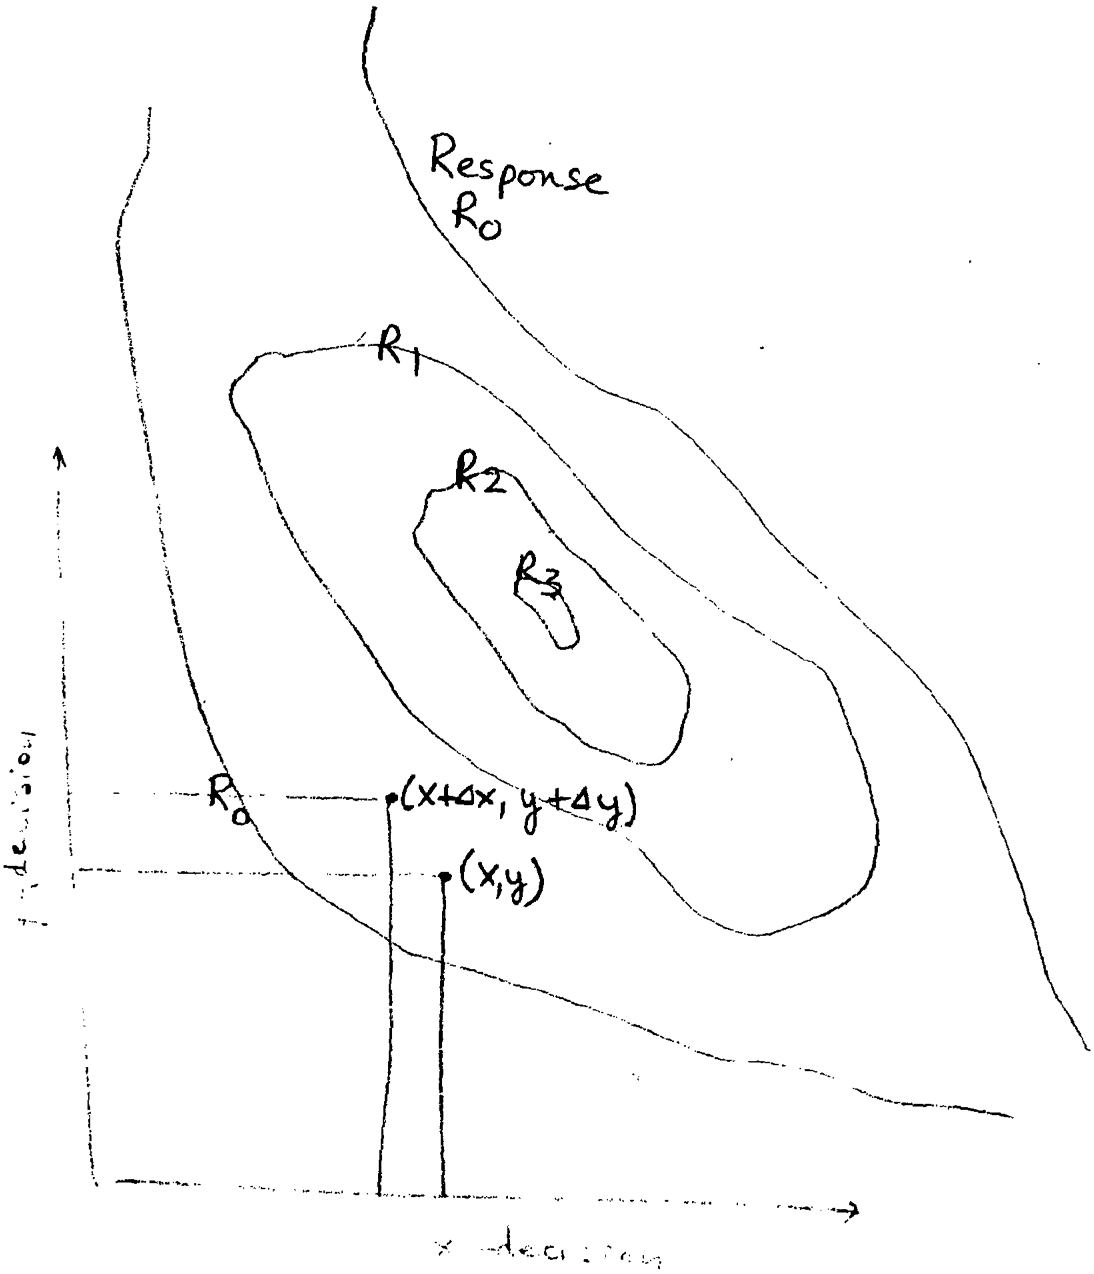

Figure 1

Table 1. System Performance Upstream-Dominate

Reservoir 1 Capacity	Reservoir 1 Target	Reservoir 2 Capacity	Reservoir 2 Target	Upstream Dominates	Downstream Dominates

Figure 2

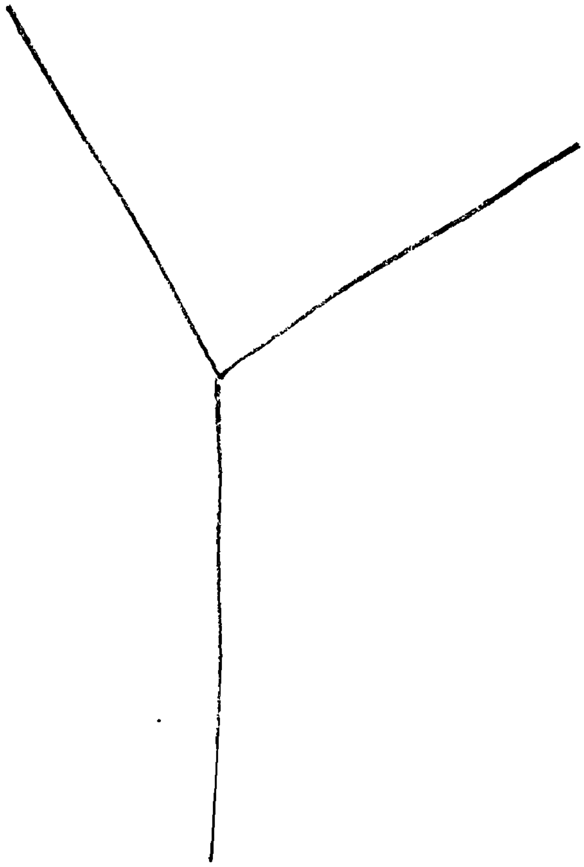

Figure 3

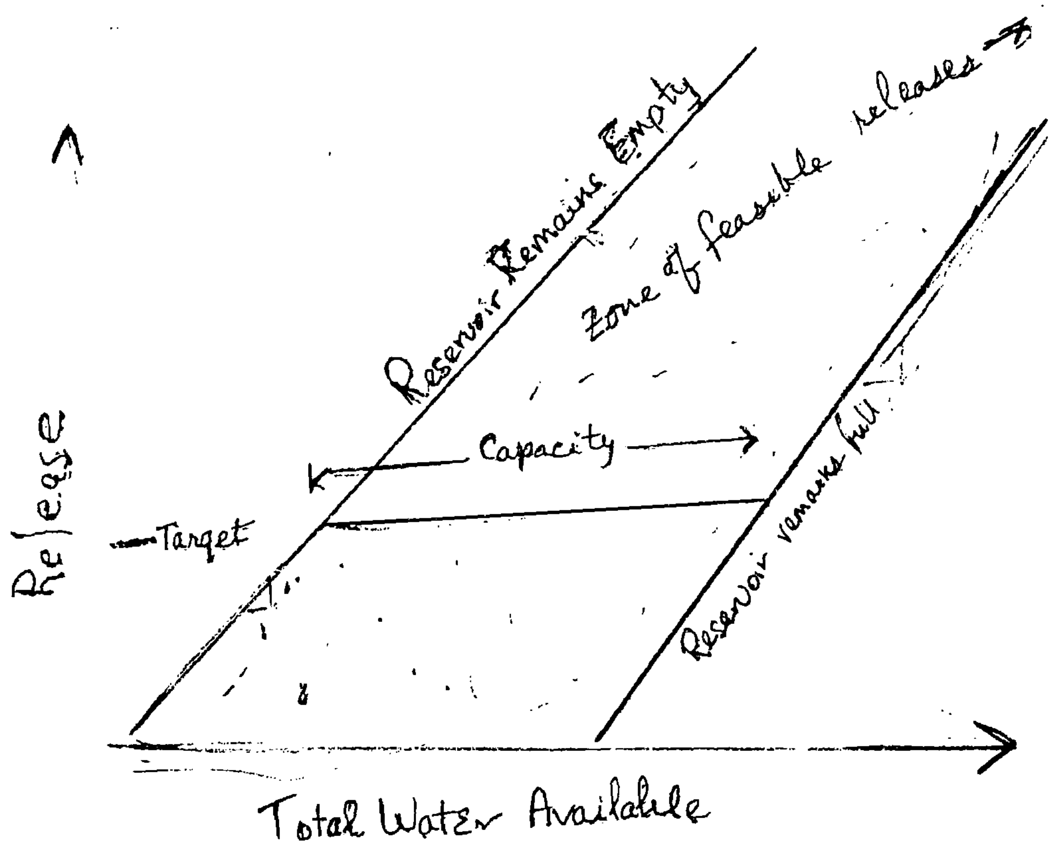

Release
Target
Reservoir Remains Empty
Zone of feasible releases
Capacity
Reservoir remains full
Total Water Available