JOURNAL OF ENVIRONMENTAL ECONOMICS AND MANAGEMENT 3, 97–112 (1976)

Environmental Policy Choice under Uncertainty[1]

W. KIP VISCUSI AND RICHARD ZECKHAUSER

Public Policy Program, John Fitzgerald Kennedy School of Government, Harvard University, Cambridge, Massachusetts 02138

Received November 25, 1975; revised February 2, 1976

The Markovian model provides an insightful structure for analyzing environmental policy decisions. This framework is applied to a variety of conceptual issues. The analysis of reswitching phenomena shows that when multiple risks are involved, there is unlikely to be an unambiguous ranking of policies according to their future-mindedness. The discussion of irreversibilities demonstrates that irreversibilities differ only in degree from other probabilistic structures; they create no special analytic complications. In the final section, the Markovian decision framework is extended to incorporate the impact of budgetary constraints on optimal decisions.

Environmental policy decisions typically involve substantial uncertainties. The long-time frame of analysis, the primitive state of knowledge about environmental implications of policies, and the intrinsically stochastic nature of many environmental occurrences have contributed to the prominence of uncertainty in recent policy debates concerning stratospheric flight, energy choices, and the trans-Alaskan pipeline, to name just a few.

It is often difficult to structure situations where uncertainty plays a critical role. This essay hopes to show that a Markovian model can frequently be helpful in this regard. The time-invariant structure for payoffs and probabilities in that model make it amenable to ready computation for specific numerical examples as well as to the development of simple frameworks within which generalizable insights can be sought. At the same time, the Markovian structure offers a richness and flexibility that enables it to capture critical elements of a variety of real world situations. Here the Markovian model is applied to each of three classes of problems that have been of interest to those with a concern for the environment.

In assessing policies that have substantial environmental implications, analysts have frequently labeled some actions as present-oriented while identifying others as consistent with a greater concern for the future. These assessments have frequently contained moral and judgmental overtones, with future orientation viewed as a virtue. Section 1 of the paper applies the Markovian model to this situation and considers whether two competing policies can be ranked unambiguously in terms of their future-mindedness within the context of a time-invariant Markov model. The

[1] Research support was received from the Analytic Methods Seminar, Kennedy School of Government, Harvard University, and through National Science Foundation grant Soc-7514258 to M.I.T., where Zeckhauser was a visitor. Alan Manne, Albert Nichols, and Edith Stokey in particular provided helpful comments. Ralph d'Arge and two referees suggested profitable directions for this reasearch.

analysis employs a traditional indicator of greater future-mindedness: A project that is inferior at discount rates above a certain level is preferred for all lower rates.

In recent years, political and intellectual developments have made us more sensitive to potential irreversibilities and to possibilities for conserving resources that are in eternally limited supply. We frequently are admonished to avoid actions that might impair welfare forever. In one of the better traditions of economics, theory has been developed to deal with reality, as exemplified by a spate of ingenious methodological investigations. Section 2 applies the Markovian model to situations where there is the possibility of irreversibilities, though there may be substantial uncertainty about the likelihood of their occurrence.

For most environmental policy choices, dollar cost is an objective that competes with environmental quality. In many circumstances, the tradeoff possibilities between the two are restricted by an imposed ceiling or constraint on expenditure levels. (Given bureaucratic realities, the ceiling frequently also serves as a floor.) Section 3 generalizes the Markovian approach to analyze situations where level of expenditure is a valued output, perhaps subject to a constraint.

Our purpose throughout is not to derive implications for policies now on the decision-making agenda, but rather to explore the more general ramifications of the methodology for environmental decisions. The principal conclusions can be demonstrated through the use of numerical examples rather than specific theorems.

1. UNCERTAINTY AND THE AMBIGUITY OF FUTURE-MINDEDNESS

In deterministic policy contexts, there may be no unambiguous way to ascertain whether one policy is more future-minded than another. Policy A may be preferred to policy B when payoff streams are discounted at rate r_1, and B preferred to A when r_2 is used. Yet policy A would once more be preferred at rate r_3, where $r_3 > r_2 > r_1$. Situations in which there is a second reversal in project preference will be referred to in this essay as instances of "reswitching."[2] For such situations, it is not possible to state which of two projects is favored as the discount rate is lowered and the future is in effect given greater weight. This ambiguity becomes important when the discount rate is endogenous, or is uncertain, or in the common instance in which policy implications are not quantified with precision but analysts nevertheless attempt to provide a simple summary of the project's temporal orientation.

Not many years ago, a harried analyst might have been tempted to dismiss the reswitching problem by observing that the choice of the discount rate should be a process quite separate from the choice of a project. The acceptable approach involved a comparison of discounted values, with the discount rate being externally specified. With a specified discount rate, the reswitching possibility becomes a curiosum, not a serous matter for policy concern.

But now matters have come full circle. Perhaps spurred by computer models of a resourceless Doomsday, the desirability of economic growth has become an open issue. Since growth rates bear a direct relationship to the marginal productivity of invested resources and since the discount rate for an equilibrium consumption stream appropriately reflects the possibilities for tradeoffs between periods (i.e., the marginal

[2] This is a narrower application than that employed by Pasinetti [23] and Samuelson [27]. Samuelson [28] many years earlier identified the possibility of multiple internal rates of return. Lorie and Savage [19], whose concern was the ambiguity of rate-of-return as a criterion for the rationing of capital, first applied this concept to the types of problems considered here.

productivity of capital in many classical models), it has become evident that the project evaluation discount rate should not be viewed as an exogenous parameter. Where the strategy choice involves projects that are large relative to the economy, the appropriate discount rate in accordance with productivity guidelines may vary depending on the strategy (i.e., projects) selected. The choice criterion "discounted expected value" will be ambiguous. Internal rate-of-return considerations may become pertinent; real-world switching problems may be an unwanted but inevitable accompaniment.

Greater concern with the environment may and should rekindle an interest in reswitching phenomena.[3] In the absence of environmental concerns, most public investment evaluations assess projects that have a period of immediate net costs followed by periods (perhaps extending to infinity) with nonnegative payoffs. If the payoff stream turns negative once again, reswitching is a possibility.[4] Such a pattern of payoffs may occur when there are environmental costs that are likely to continue far into the future.[5]

In view of the increased importance of uncertainty, particularly with respect to environmental problems, it seems only fitting that the reswitching discussion should be extended to contexts where uncertainty plays a major role. In a trivial sense, any deterministic reswitching demonstration may be viewed as a special case of the more general world of uncertainty, with all probabilities either 1 or 0. To convert such demonstrations to nondegenerate examples, we merely need to substitute very low probability outcomes offering slightly different payoffs.

By applying a bit more ingenuity to the same type of demonstration, we can develop more significant demonstrations of reswitching under uncertainty. Consider four deterministic projects which are associated with the payoff vectors: $A_1 = (9, 21, 0.01)$, $A_2 = (0, 0, 34)$, $B_1 = (0, 11, 20)$, and $B_2 = (0, 29, 4)$, where the ith entry is the payoff in period i. These payoff vectors have the properties:

(1) At interest rates above 3.4%, A_1 is preferred to B_1; otherwise B_1 is preferred.
(2) At interest rates below 3.5%, A_2 is preferred to B_2; otherwise B_2 is preferred.

Introduction of uncertainty may introduce reswitching in the choice of optimal policies, even though only single switches occurred in the deterministic cases. For example, the lottery $(0.5A_1; 0.5A_2)$ is preferred to the lottery $(0.5B_1; 0.5B_2)$ at interest rates below 1.12% and above 9.99%.

This type of analysis could be elaborated in a traditional decision tree format. Each controlled decision or random event would lead not to a probability distribution over a single payoff, but rather to a distribution over time vectors of payoffs. Employing the usual procedures of backward induction to optimize, one might discover at a

[3] Albin [2] remarks on the reintroduction of horses and the phasing out of mechanical hauling processes in logging operations in the Pacific Northwest. He suggests that this reswitching in technology (not to be confused with future-mindedness reswitching), may be due to enhanced concern with the environment in particular to a desire to protect the floor of the forest.

[4] With competing projects, traditionally shaped streams of returns no longer offer guarantees against reswitching. The stream of differences between the period payoffs of the projects is crucial. These differences can well change sign twice, even though stringent regularity conditions are imposed on the streams of the individual projects. (For example, it can be required that the first and second time derivatives of the individual project streams never change sign.) The presence of two such changes in the difference stream is sufficient to permit the possibility of reswitching.

[5] A numerical example involving energy policy choices in which environmental costs create a reswitching situation is provided by Viscusi and Zeckhauser [32].

particular binary-choice node that one branch was superior for high and low interest rates, but not for intermediate rates. The decision tree approach may be extended to develop the relationship of reswitching to such concepts as the expected value of perfect information.

Markovian Models of Policy Decision

In part because of the computational unwieldiness of multiperiod decision trees, this paper focuses on the Markovian model as a tool for analyzing uncertainty situations. A Markovian model is described by a matrix, called a Markov matrix or system, whose entries are transition probabilities. The entries indicate the likelihood that if the system is in a particular state now it will be found in a specified state a period hence. For the examples in this paper, the states represent different levels of environmental quality. The transition probabilities and the payoffs for each state depend upon the environmental protection policy that is selected, but they do not depend upon time. When the entries in the matrix are zeros and ones, the model represents the techniques of optimal control theory and dynamic programming that have of late been profitably applied to problems of the environment.[6] The interest here is more in stochastic situations, those where many of the matrix entries lie between zero and one.

The Markov model achieves its greatest expository advantage when the alternatives for choice remain identical over time, and when the number of periods is large. Discrete interpretations of continuous-outcome situations are frequently instructive, and Markov systems can be applied to situations where states, policies, and time are continuous. It is only when the range of outcomes is infinitely large that the Markov representation is clearly inappropriate.

Note also that the Markov formulation can be applied to situations where many independent units are circulating among different states of the world. For example, these might be physical units of polluting substances, perhaps transiting between a storage facility, the open environment, and our food. Alternatively, the Markovian analysis can be applied to situations where the "transition probabilities" really represent the proportion of some substance that will transit from one state to another. Thus, for example, a model of technologies to improve water quality might have an aeration plant that removed $X\%$ per period of a particular pollutant that is dumped into a stream. Of course, if the number of units facing independent trials is large, observed proportions will closely approximate per-unit probabilities.

Markov systems for policy decisions have probably received their most widespread attention in environmental economics in application to reservoir and river-basin management.[7] The uncertain variable considered in these contexts is water inflow. Demand magnitudes (which are, in effect, surrogates for benefits) are also sometimes treated as uncertain.

[6] The importance of the Markov assumption in optimal control models is emphasized by Pontryagin *et al.* [24]. Dynamic programming methodologies likewise assume that this separability property holds. As Bellman [6] notes: "The basic property that we wish F (the reward function) to possess is one of Markovian nature, to wit: after any number of decisions, say k, we wish the effect of the remaining $N-k$ stages of the decision process upon the total return to depend only upon the state of the system at the end of the kth decision and the subsequent decisions."

[7] The papers by Dorfman [9] and by Thomas and Watermeyer [31] followed this approach in their seminal contributions to environmental economics and management.

Markovian decision models should be more extensively exploited in the examination of environmental situations.[8] Problems of environmental control and regulation frequently combine the critical elements of stochastic outcomes and closed-loop structure. The first element means that we do not know what will happen. The second implies that certain situations will recur (say high levels of air pollution) and that whatever it is best to do the first time a situation is met (for example, generate more power at Plant A and less at Plant B) will also be optimal next time around. Such problems are readily analyzed with the aid of Markov transition matrices.[9]

It might be expected that in a time-invariant model such as this, there will be no ambiguity as to which of two policies was more future-minded. The examples below indicate otherwise, however.

A Pollution-Control Model with Uncertainty and Reswitching

To demonstrate possibilities for reswitching under conditions of uncertainty, we employ a simplified example from controlling radioactive wastes. The Nuclear Regulatory Commission is formulating a policy for the disposal of radiation waste from nuclear plants. Strategy A would bury the waste in empty salt mines. Strategy B would store them in a ground-level facility. There is less risk of contamination from strategy A, assuming that there are not present leaks in the salt mines. (The AEC's unfortunate experience at Lyons, Kansas (where an attempt was made to store radioactive wastes in caverns that had been pierced by oil and gas wells) suggests there is a real danger of undetected leaks.) But if some contamination should be detected, strategy B will offer greater resiliency. That is, we may be able to recover contaminating wastes if they are being kept at ground-level, or at least some portion of them, and prevent them from doing further harm.

For simplicity, the model is formulated as if we could never recover from contamination from wastes stored in the salt mines. Radioactive materials decay, of course, and that will help diminish damages over time. This decay phenomenon can be accommodated in the analysis. Say the rate of decay of the waste is x per period and r is the interest rate. Then the discount factor that should be applied to payoffs is $(1 - x)/(1 + r)$, for that factor gives the per period rate of decline in damages from a contamination incident. Other forms of recovery of wastes, of course, will reduce damage amounts in direct proportion to the quantities recovered.

State 1 plays an initial sorting role for strategy A and never enters the valuation calculations. It determines whether the salt mines are or are not initially sealed to the outside. If they are sealed initially, a 0.98 probability, then there is only a 0.001 chance per period that a geological disturbance such as an earthquake will lead to outside contamination.

There is a 1% chance each period that radiation wastes in ground-level storage will lead to contamination. If they do, then there is a 20% chance of recovery each period. (Alternatively, this could represent a situation in which 20% of the initially contaminating material was recovered in a period.)[10]

[8] Howard [15, 16] developed the methodology closest to that employed here.

[9] As we demonstrate with an irreversibility example below, Markov matrices can also be employed when elements of learning are involved, that is when transition probabilities may be subject to updating.

[10] In some of the examples below we explicitly allow for learning. Learning might be helpful in developing a strategy for storing radiation wastes, assuming that some of the transition probabilities in the matrices were really random variables subject to estimate. In this example, we might learn for instance about the probability of recovery from the contamination state.

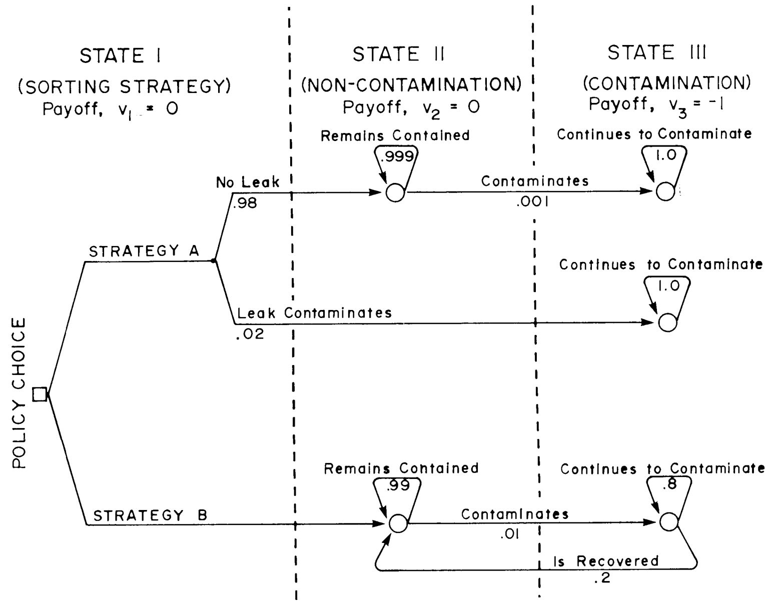

FIG. 1. Nuclear waste policy choice.

The wastes cause no damage, i.e., offer a payoff of 0, in a noncontaminating state. In a period in which there is contamination, the payoff is arbitrarily scaled to be -1. Figure 1 illustrates the Commission's choice problem. The objective is to select the strategy that maximizes the expected discounted rewards. It is to choose the strategy that yields the

$$\operatorname*{Max}_{S} \sum_{k=0} \sum_{i} v_i p(i \mid S, k)/(1+r)^k,$$

where $p(i \mid S, k)$ gives the probability that the nuclear wastes are in state i in period k if strategy S is employed, and v_i indicates the payoff for state i. (This analysis is presented for a single injection of wastes. It may be extended to repeated injections.)

For discount rates below 5.27% or above 18.61%, strategy B is superior.[11] For discount rates in the interval bounded by these switch points, strategy A is superior. Thus reswitching is encountered in a time-invariant Markov process. Careful consideration of Fig. 1 makes it evident that reswitching might occur. Consider the extremes. At extraordinarily high discount rates only initial period damage matters; strategy A is inferior. For zero discount rates, long-term residency probabilities are what is of consequence. In the long run if strategy A is employed, the wastes remain forever in state 3, the most damaging state. Because it keeps the probabilities of passage from state 2 to state 3 relatively small, strategy A is attractive for intermediate discount rates. Figure 2 plots the difference in expected discounted damage for the two strategies as a function of the discount rate.

This example illustrates that a strategy leading to unattractive trapping states is potentially undesirable. However, the presence of an absorbing state is not a necessary

[11] Howard [16] presents a methodology for solving this problem. The key algorithm used is $u = (I - \beta P)^{-1}q$, where u = n-dimensional vector with u_i the expected discounted payoff starting in state i; $I = n \times n$ identity matrix; β = discount factor, $1/(1 + r)$; $P = n \times n$ transition matrix; and q = n-dimensional vector with q_i the expected immediate rewards starting in state i.

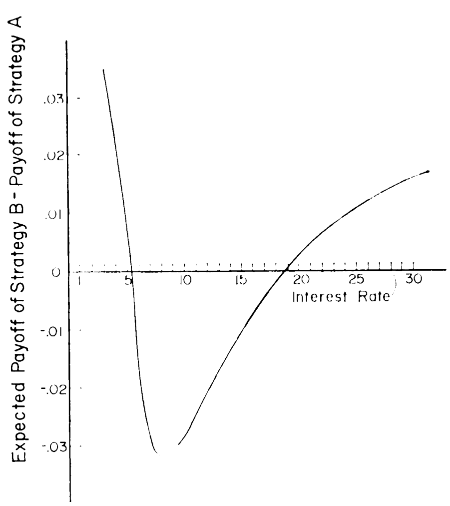

FIG. 2. Expected discounted cost comparison.

condition for reswitching. If, for example, strategy *A* provided a very small probability of exiting from the contamination state, there would be no absorbing state in the problem. Yet, the relative cost patterns of the strategies and the switch points would not be affected significantly.[12]

The implication of this latter observation foreshadows our later discussion of irreversibilities. The presence of irreversibilities need not create major discontinuities in the attractiveness of policies, for minor variations in transition probabilities will alter the expected discounted payoffs of policies only slightly. The dramatic impact attributed to irreversibilities in some of the literature stems in part from the use of deterministic models. These models investigate situations in which the probability of remaining in an attractive state is either zero or one. It is not surprising that policy payoffs are sharply affected, with conclusions such as, "If we continue to use aerosol cans, the ozone layer will be irreparably depleted." Most real world situations are less stark.

Those who are alerting us to irreversibilities may believe that a nonnegligible probability that a policy will lead to an irreversible outcome is being ignored; in effect it is being treated as if it were zero. Some of them carry their analysis to the other extreme and assume that the irreversibility is certain. This seems to be the argument, for example, of some critics who are warning society about the time-bomb nature of accumulations of nuclear wastes. It is important to recognize, however, that variations along a probabilistic continuum produce no discontinuous changes in valuation.[13] In the example just cited, therefore, it is important to assess whether the danger of widespread exposure to radiation in a decade is 0.01, 1, or 10%. Our choice of policies for energy generation might be very different for the three circumstances.

[12] To verify this assertion, suppose that once strategy *A* enters state 3, there is a 0.001 probability that state 1 will be entered and a 0.999 probability that the process will remain in state 3. Even though state 3 is no longer absorbing, this modification increases the attractiveness of strategy *A* only slightly. Strategy *B* is now superior if the interest rate is below 18.90 or above 4.92% (figures that differ little from our earlier switch points of 18.61 and 5.27%).

[13] This result may no longer hold if a zero discount rate or some nondiscounting criterion is employed to assess eternal welfare. In that case, even a small probability of entering a state that will be forever inferior may render unacceptable a strategy that risks the irreversibility.

Reswitching Policy Vectors and Policies for Particular States

Markovian decision problems involve two broad types of policy choices: We may choose between entire strategy vectors (i.e., arrays prescribing a policy for each state) or we may be able to select the policy for each state individually. In the pollution example, the choice was between vectors of policies. Here we will provide examples that we believe have the minimum number of states required if reswitching is to be a possibility.[14]

Consider the choice between alternative strategy vectors C and D. The Markov process is assumed to begin in state 1 for each, and in each state the cost to use either strategy is the same. Let

$$C = \begin{bmatrix} 0.5 & 0.5 & 0 \\ 0 & 0.98 & 0.02 \\ 0 & 0 & 1 \end{bmatrix} \quad \text{and} \quad D = \begin{bmatrix} 0 & 1 & 0 \\ 0.3 & 0.695 & 0.005 \\ 0 & 0 & 1 \end{bmatrix},$$

where the ijth element of each matrix is the transition probability from state i to state j. Let the payoff vector v be given by

$$v = [0,\ 1000,\ -100],$$

where v_i is the reward associated with state i.

The choice problem is to select the strategy whose transition matrix provides the greatest expected discounted payoff. Each policy leads eventually to state 3, an undesirable trapping state, i.e., a state for which the probability of exit is zero. Strategy C will on average reach that state sooner. For the parameters of this example, strategy D is preferred for interest rates below 11.4% and above 15.0%, while strategy C is preferred in the intermediate range. If there are at least three states, reswitching may occur when choosing between alternative vectors of policies.

Often one cannot choose between alternative vectors of policies, but can only vary the policy for a particular state.[15] For example, suppose the choice is between two alternative transition matrices that are identical except for a single row. The row might represent a choice of policy for a particular environmental context. We believe that for reswitching to occur in this situation, there must be at least four states.

Consider a numerical example where $v = [50, 200, 0, 200]$. The transition matrix is defined for transition from states 2, 3, and 4 as shown below.

$$\begin{bmatrix} \multicolumn{4}{c}{\text{TO BE SELECTED}} \\ 0 & 0.5 & 0.5 & 0 \\ 0 & 0 & 0.99 & 0.01 \\ 0 & 0 & 0 & 1.0 \end{bmatrix}.$$

The decision problem is to decide between two alternate strategies for transiting out of state 1, i.e., to choose between two alternative top rows. They are $1_a = (1, 0, 0, 0)$ and $1_b = (0, 1, 0, 0)$. The process commences in state 1.

For discount rates between 4.5 and 11.2%, 1_a is superior to 1_b. For very low discount rates (below 4.5%) the fact that 1_b ultimately leads to the attractive, trapping state 4 makes that option superior. However, for rates between 4.5 and 11.2%, the fact that 1_b leads us to reside for a while in the unattractive, but nontrapping state 3

[14] These limiting cases are put forward as conjectures. The minimum requirements for reswitching to occur have yet to be derived in rigorous fashion.

[15] This is actually a restricted case of the problem involving a choice between policy vectors.

outweighs that option's eventual superiority. Reswitching is observed once again.

These results with uncertainty models reinforce the earlier findings in which environmental concerns were responsible for reswitching. In a great variety of instances, there can be no assurance that one project will appear unambiguously more future-minded than another.

2. APPLICATIONS TO IRREVERSIBILITIES

A series of recent papers has considered the conceptual issues posed by environmental irreversibilities.[16] These irreversibilities make it impossible to achieve certain environmental outcomes, or indeed to make specific policy choices once particular decisions have been made. In short, irreversibilities involve situations where the implications and consequences of past policies cannot be terminated at will. This observation flows naturally from the use of a Markovian decision framework, for it emphasizes the probabilistic dependence of future outcomes on present decisions and the immediate outcomes they produce.

We will define a potential irreversibility as the existence of a trapping state or set of trapping states within the Markov transition matrix. Policy concern with respect to the irreversibility will revolve around such questions as: How likely is it that the trapping states will be entered? How long on average will it take to transit to the trapping states? What expected discounted payoffs are associated with the trapping states? Intelligent policy formulation must address these questions.

Several insights into irreversibility can be afforded by a Markov model. First, by highlighting the contrast with states that can be reentered, it offers a powerful illustration of the potential impact or likely lack of impact of an irreversibility.

Second, the Markov approach demonstrates that states that are not unattractive in themselves should perhaps be avoided even at considerable cost because they increase the likelihood of future transit to undesirable states. This lesson becomes obvious once the Markovian model has been formulated, and may be important in environmental problems with potential future irreversibilities. Choices on technologies for electricity generation or the regulation of aerosols or stratospheric flight fall within this category.

Third, and most important, many of the more interesting unresolved questions relating to irreversibility involve uncertainty, and uncertainty is the driving force of the Markov model.

How shall we deal with situations where an irreversibility is a risk but hardly a certainty? We can distinguish two quite different types of uncertainty regarding irreversibilities. The irreversibility of a particular state (or group of states) may be certain, but the likelihood that a particular policy will lead to the irreversible situation may not be. (Say we adopt an unrestrictive radiation-control strategy. Individuals will surely be exposed to long-lived nucleides. In the future, however, we may have procedures to arrest the development of the malignancies these nucleides would otherwise induce. The generation of substantial levels of fatal cancers is not a certain irreversible consequence of a lax radiation control strategy.) Alternatively, it may not be known with certainty whether a particular state or group of states is irreversible. For example, we do not know to what extent depletion of the ozone layer is an

[16] See Refs. [1, 3–5, 7, 8, 10–13, 33]. As one might expect from the volume of papers, a consensus concerning the appropriate definition and resolution of the issues posed by irreversibility has yet to emerge.

irreversible process. Each of these classifications of irreversibility can be made only with respect to the feasible policy options. A state may be irreversible for one policy, but not for another which permits transitions from the formerly irreversible state. And it may be uncertain whether in the future some option will become available which will make an "irreversibility" in fact reversible.

We provide an example of a situation where it is not known whether a particular state is irreversible. In a three-state world assume that there is a 50% chance that unattractive state 3 is a trapping state, and hence represents an irreversibility. The policy options are E and F; they involve only the structure of transition probabilities out of the first two states:

$$E = \begin{bmatrix} 0.4 & 0.6 & 0 \\ 0.6 & 0.35 & 0.05 \\ \multicolumn{3}{c}{\text{PREDETERMINED}} \end{bmatrix} \quad \text{and} \quad F = \begin{bmatrix} 0.88 & 0 & 0.12 \\ 0 & 1 & 0 \\ \multicolumn{3}{c}{\text{PREDETERMINED}} \end{bmatrix}.$$

The PREDETERMINED third row of the transition matrix is equally likely to be (0.5, 0, 0.5) or (0, 0, 1). The payoff vector is (10, 5, 0); the discount rate is 10%. If this were a one-time decision, we could readily compute the expected discounted value for each of the two options and fold back the decision tree, as illustrated in Fig. 3. The discounted expected value for E is 72.9; for F it is lower at 66.9. In a one-shot trial, strategy E, which incurs a lower risk of the irreversible state, is preferred.

Matters change, perhaps counterintuitively, if learning is permitted. Suppose that we cannot observe the structure of row 3, but we can observe whether we transit out of that state. Each sequential strategy has its own transition matrix. Strategy EE corresponds to the exclusive use of strategy E regardless of the outcome of any transitions from the potentially irreversible state. Strategy FF is defined analogously. These two strategies were considered above using the decision tree framework. (The transition matrices EE and FF require a 7×7 format.)

Strategy EF utilizes strategy E until a transition from the potentially irreversible state is observed to occur. At that time (which will never come if state 3 proves to be irreversible), strategy F becomes the choice. The transition matrix for strategy EF is rather complex, involving the 10×10 display presented in Table I. It incorporates three submatrices. The first, made up of rows 2, 3, and 4, represents the use of strategy E when state 3 is irreversible. The second submatrix, involving rows 5, 6, and 7, describes transitions when strategy E is employed and state 3 is reversible. The third submatrix, defined by rows 8, 9, and 10, shows the use of strategy F after the reversi-

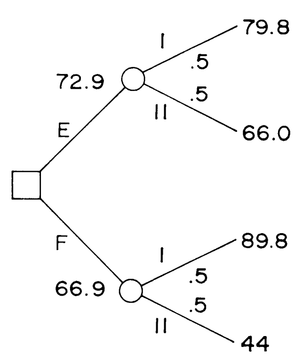

FIG. 3. One-time only choice.

TABLE I

Transition Matrix for Strategy *EF*

Key to matrix											
First period chance move out of State 1		0	0.2	0.3	0	0.2	0.3	0	0	0	0
Strategy *E*, Irreversible world	State 1	0	0.4	0.6	0	0	0	0	0	0	0
	State 2	0	0.6	0.35	0.05	0	0	0	0	0	0
	State 3	0	0	0	1	0	0	0	0	0	0
Strategy *E*, Reversible world	State 1	0	0	0	0	0.4	0.6	0	0	0	0
	State 2	0	0	0	0	0.6	0.35	0.05	0	0	0
	State 3	0	0	0	0	0	0	0.5	0.5	0	0
Strategy *F*, Reversible world	State 1	0	0	0	0	0	0	0	0.88	0	0.12
	State 2	0	0	0	0	0	0	0	0	1	0
	State 3	0	0	0	0	0	0	0	0.5	0	0.5

bility of state 3 has been observed. Each of states 1, 2, and 3 is represented three times in the matrix. State 1 appears in rows 2, 5, and 8; state 2 represents rows 3, 6, and 9; state 3 is rows 4, 7, and 10. The payoff vector obviously repeats itself in triplets; it is

$$v = [10, 10, 5, 0, 10, 5, 0, 10, 5, 0].$$

(Note that there is no way to get to the first row; the first element of v is arbitrary.)

The first row of the matrix for *EF* gets the process going with a chance move out of state 1, the starting state. It is 40% likely that the process will be in state 1 after the first move. Since state 1 for strategy *E* is represented by the first rows of two equally likely submatrices, the entries in columns 2 and 5 are both 0.20. A similar process produces the 0.30 in columns 3 and 6. Inspection of row 7 of the 10×10 matrix shows what happens when it is discovered that state 3 is reversible and a switch is made to strategy *F*.

For the values in this problem, sequential strategy *EF* is optimal, having a discounted expected value of 73.6, as opposed to *EE*'s 72.9.[17] The appropriate policy, therefore, is to begin with strategy *E* and switch to the formerly more risky strategy *F* when and if it is learned that there is no absorbing state. The sequential policy for this problem is superior to policies that fail to exploit the information generated by the process.

A general lesson can be found here. Although an action that does not risk irreversibilities may offer a higher expected value than any other nonsequential strategy, it may at the same time be inferior to sequential strategies that enable us to learn whether the irreversibility does in fact exist (or at least more about its likelihood) before committing ourselves for the future. The possibility of acquiring information had no effect on our initial policy in the example in the text. Other examples may lead us to be more cautious in some instances or to take greater initial risks in others. The nature of the optimal policy depends on the structure of the problem: the transition probabilities, the payoff vectors, the opportunities for learning, and the ability to alter a policy in subsequent periods. The presence of irreversibilities causes no qualitative change in the nature of the problem or in the policies that should be selected.[18]

[17] For lower discount rates, the difference between *EE* and *EF* increases. For a discount rate of 4%, the difference is 3.9.

[18] Much of the confusion with respect to the impact of irreversibilities can be traced to a misunderstanding of the concept of option values. See [35] for a discussion that is pertinent to this issue even though it does not deal with irreversibilities explicitly.

Where there is uncertainty, there may be learning. Thus, in a variety of policy contexts there can be adaptive responses to increased knowledge about a state's irreversibility. If we do not know whether an endangered bird species can survive a significant alteration in its present habitat, we may not need to choose right now between the full alteration and none at all. Monitoring the birds' responses to incremental changes may tell us whether or not the birds are likely to survive the significant alterations under consideration. Learning about irreversibilities should be more than a process of intellectual model building. When confronted with real policy choices, we should attempt to learn about the values of the parameters that describe the likelihood that the potential irreversibilities will occur, and the consequences when they do.[19]

3. UNCERTAINTY AND POLICY CHOICE WITH COST AND ENVIRONMENTAL QUALITY AS VALUED ATTRIBUTES

The traditional Markovian decision model details procedures for choosing the appropriate strategy to maximize a discounted expected payoff, when the choice among discrete strategy options affects both transition probabilities and period payoffs. A third dimension may be of consequence as well; it is the dollar expenditures that will be required (on a discounted, expected value basis) to pursue the alternative strategies. For many situations, dollar expenditures can be employed as surrogates for the strategies as they apply to particular states of the world. We may be interested, for example, in how much we should spend improving environmental quality in each of a variety of possible circumstances, where it is well understood how different levels of budget would be allocated. To deal with situations where money as well as environmental quality play a role, we extend our model in two ways. First, we allow for policy options that are represented by continuous variables. Second, we will assign a budgetary cost to each state-option pair, and make this cost a consideration in the final policy choice.

Thus discounted expected dollar costs, C, will enter the objective function as an argument along with discounted expected environmental quality. Environmental quality is scaled in benefit units, b, measured from some reference standard. B indicates discounted expected benefits. The objective function, W, can then be written

$$W = f(B, C).$$

Note that this formulation implies that the preference ordering among alternative streams of environmental benefits will not depend in any way on the stream of dollar expenditures. In some circumstances, the objective function will take the form

$$f(B, C) = B - \lambda C.$$

More generally, for any form of the f function, given B and C it will be possible to define the marginal tradeoff rate between the two valued attributes. That assumption will be important for our formulation below.[20] This formulation implies that the rate

[19] Ralph d'Arge has described a research project which would take a historical look at perceived potential irreversibilities from the past. A careful retrospective investigation would reveal whether our society is too cautious or too bold when confronting potential irreversibilities.

[20] Note we deal only with expected values here. This approach is more general than it may appear. The period b_i's may already be scaled to show that period "payoff" is not linear with physical values. The total variability of discounted sums may be small if the number of periods is large and the probabilities in the Markov matrix allow for frequent transit among the states. Finally, the size of the dollar

at which the decisionmaker is willing to trade dollar expenditures for environmental quality will be the same in every period.

A particular version of this formulation would place a constraint on B or C. Thus the objective might be to maximize discounted environmental quality subject to some expected level of expenditure. Alternatively, the problem may be to find the cheapest way to achieve some required expected level of environmental quality.[21]

The policy choice problem is now straightforward. For each state, compute a simulated payoff

$$v_i = b_i - \lambda c_i,$$

where b_i and c_i are the environmental benefits and budgetary expenditures, respectively. Either the b_i's or the c_i's or both may be a function of the selected policy. For example, different policies may entail different levels of expenditure within particular states, implying that the c_i's will be policy specific. Similarly, the policies may affect the within state benefits, not only the probabilities of transition among the states. If the value of λ is known, conventional optimization techniques can be used to determine the optimal policy vector when state payoffs are defined in the manner described above.

When λ is not known explicitly, one can employ the following iterative process. For a particular value of λ, determine the optimal policy vector. Compute the values of B and C associated with that vector and substitute them into the objective function. By systematically varying λ, the frontier of production possibilities is traced for B and C. Usually, it will only be necessary to compute a small portion of the frontier in the neighborhood of the optimum, where it is tangent to one of the indifference curves implied by the objective function.

If the problem is instead one of maximizing discounted expected environmental payoffs subject to a constraint on discounted expected expenditures, λ is increased (if C is too high) or decreased (if C is too low) until the optimal policy vector satisfies the budgetary constraint. A simple example will illustrate this optimization process.

A government agency is regulating the use of a pesticide. The accumulation of the pesticide within the environment depends on a number of uncontrollable factors; it also depends on the amount of the pesticide that people employ. For simplicity, let us say that there are two levels of accumulation for the pesticide at the beginning of the period, high or low. The policy decision is how strictly to regulate pesticide use, given the measured accumulation at the beginning of the period. Each particular level of regulation, x, can be uniquely associated with its dollar cost in terms of production foregone or alternative pest-elimination costs incurred because certain uses of the pesticide were prohibited. This suggests that the policy choice, x, can be uniquely correlated with a value of c.

---

expenditure on this area of environmental control may be small relative to other expenditures out of the same budget, in which case variations around expected expenditures will not be of consequence for the evaluation of welfare.

[21] In a vast range of regulatory situations, environmental quality is highly uncertain. Frequently standards are employed in such situations, so that dollars are not traded directly against environmental quality. The least efficient mode of standard setting states that quality can never fall below some specified level. The next higher level of sophistication would specify a maximum percentage of the time during which some quality minimum could be violated. A system of much greater attractiveness, of course, would assign a value to all outcomes and specify a minimum aggregate score over time. Better still, it would suggest that that score had only to be met on an expected value basis.

The transition matrix is given by

		Next period	
		Low level	High level
This period	Low level	$p_1(c_1)$	$1 - p_1(c_1)$
	High level	$p_2(c_2)$	$1 - p_2(c_2)$.

We would expect $p_i(c_i)$ to be a positive function of c_i; that is, the more stringently pesticides are regulated, the more likely a low level of accumulation will be observed next period.

The environmental payoff in a period depends both upon the initial measured cumulation of the pesticide and the amount of spraying during that period. We can represent these environmental payoffs as $b_1(c_1)$ and $b_2(c_2)$ for the initial low and high levels, respectively. We would expect these functions to be positively related to their arguments. The vector of hypothetical payoffs, v, is given by

$$v_1 = b_1(c_1) - \lambda c_1$$

and

$$v_2 = b_2(c_2) - \lambda c_2.$$

For simplicity, assume that the process starts in state 1, and that u_1 is the present value associated with the process. For each particular value of λ, we want to maximize u_1. The per period discount factor is $\beta = 1/(1 + r)$, and the time horizon is infinite. The first-order maximization conditions for c_1 are

$$\frac{\partial u_1}{\partial c_1} = 0 = (\beta - \beta^2)\frac{\partial p_1}{\partial c_1}[(b_1 - \lambda c_1)(p_1 - \beta p_1 + \beta p_2 - \beta p_1 p_2)$$
$$+ (b_2 - \lambda c_2)(1 - p_1 - \beta p_2 + \beta p_1 p_2)] + [1 - \beta + (p_1 - p_2)(-\beta + \beta^2)]$$
$$\times \left\{ \frac{\partial p_1}{\partial c_1}[(b_1 - \lambda c_1)(1 - \beta) - (b_2 - \lambda c_2)] + \left(\frac{\partial b_1}{\partial c_1} - \lambda\right)(p_1 - \beta p_1 + \beta p_2)\right\}.$$

The analogous condition for c_2 is a similarly complicated expression. Solution of the two first-order conditions for c_1 and c_2 provides the optimal policy choice vector.

Consider a concrete example. Suppose the discount factor is 0.9 and that the functions take the following form.

$$p_1(c_1) = 1 - e^{-c_1}, \qquad p_2(c_2) = 1 - e^{-0.8c_2},$$
$$b_1(c_1) = 2 - e^{-0.1c_1}, \qquad b_2(c_2) = 1 - e^{-0.5c_2}.$$

If λ is equal to 0.5, the optimal policy expenditure vector is (0.7, 1.2), which produces a discounted expected payoff of 3.27. Alternatively, suppose that the budgetary shadow price is not given, but that the budgetary constraint is known to equal 14.0. The optimal policy choice is (1.2, 1.9), which produces a discounted value of 5.54 and is associated with an implicit budgetary shadow price of 0.3. More generally, Fig. 4 plots the opportunities frontier available. It was derived by maximizing u_1 for alternative values of λ, and plotting the associated maximizing values of B and C. Discounted expected benefits are a concave function of discounted expected budgetary costs. The appropriate policy choice is that given by the point of tangency with an indifference curve implied by the objective function.

The Markov model in this situation enabled us to model a complex environmental decision problem in a streamlined fashion. Costs of environmental regulation, the

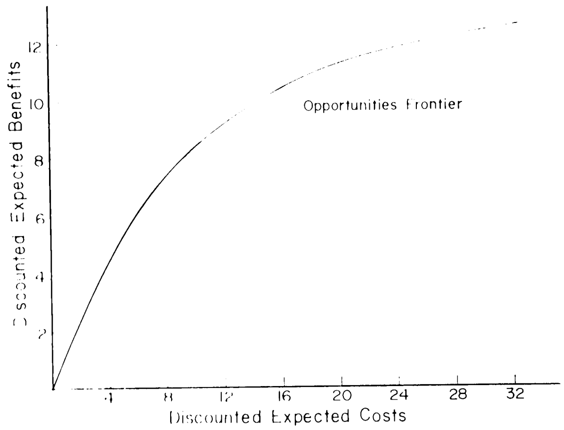

FIG. 4. Opportunities frontier for policies of pesticide control.

short-run benefits it confers, and its long-run explicitly uncertain consequences were all introduced into the policy choice process. It appears that no equivalently simple model offers this richness. On a benefit/cost basis, Markovian models of environmental decisions are most attractive.

4. CONCLUSION

Those with a concern for resource limitations, the environment, and the possible irreversibility of policy actions should profit from the systematic inclusion of uncertainty in their analyses. A Markovian model is frequently useful here. As we have shown, it can lend insights into analytic topics of policy concern such as reswitching and irreversibilities. Moreover, it can help with the recurring problem of formulating intelligent strategies for environmental quality when dollar costs and short- and long-run environmental quality compete as objectives for policy choice.

REFERENCES

1. A. E. Abrassart and D. D. McFarlane, The economics of environmental preservation: Comment, *Amer. Econ. Review* **64**, 1025–1029 (1974).
2. P. S. Albin, Reswitching: An empirical observation, a theoretical note, and an environmental conjecture, Working Paper No. 74-09, New York University (1974).
3. K. J. Arrow, Optimal capital policy with irreversible investment, *in* "Value, Capital, and Growth," (J. Wolfe, Ed.), Aldine, Chicago, Ill. (1968).
4. K. J. Arrow and A. C. Fisher, Preservation, uncertainty, and irreversibility, *Quart. J. Econ.* **88**, 312–319 (1974).
5. K. J. Arrow and M. Kurz, Optimal growth with irreversible investment in a Ramsey model, *Econometrica* **38**, 331–344 (1970).
6. R. Bellman, "Adaptive Control Processes: A Guided Tour," Princeton Univ. Press, Princeton, N.J. (1961).
7. M. L. Cropper, The regulation of activities with potentially disastrous consequences, Working Paper, Univ. of California, Riverside, Ca. (1974).
8. R. Cummings and V. Norton, The economics of environmental preservation: Comment, *Amer. Econ. Review* **64**, 1021–1024 (1974).
9. R. Dorfman, Mathematical models: The multistructure approach, *in* "Design of Water-Resources Systems," (A. Maass *et al.*, Eds.), Harvard, Cambridge, Mass. (1962).
10. A. C. Fisher and J. V. Krutilla, Valuing long-run ecological consequences and irreversibilities, *J. Environ. Econ. Manag.* **1**, 96–108 (1974).

11. A. C. Fisher, J. V. Krutilla, and C. J. Cicchetti, The economics of environmental preservation: A theoretical and empirical analysis, *Amer. Econ. Review* **62**, 605–619 (1972).
12. A. C. Fisher, J. V. Krutilla, and C. J. Cicchetti, The economics of environmental preservation: Further discussion, *Amer. Econ. Review* **64**, 1030–1039 (1974).
13. C. Henry, Investment decisions under uncertainty: The 'irreversibility' effect, *Amer. Econ. Review* **64**, 1006–1012 (1974).
14. J. Hirshleifer, "Investment, Interest and Capital," Prentice–Hall, Englewood Cliffs, N.J. (1970).
15. R. A. Howard, "Dynamic Probabilistic Systems," Vols. I and II, Wiley, New York (1971).
16. R. A. Howard, "Dynamic Programming and Markov Processes," MIT Press, Cambridge, Mass. (1960).
17. F. Knight, "Risk, Uncertainty, and Profit," Houghton Mifflin, Boston, Mass. (1921).
18. K. Lancaster, A new approach to consumer theory, *J. Polit. Econ.* **74**, 132–157 (1966).
19. J. Lorie and L. J. Savage, Three problems in capital rationing, *J. Business* **28**, 229–239 (1955).
20. A. S. Manne, Waiting for the breeder, *Review Econ. Studies, Symp.* 47–66 (1974).
21. A. S. Manne and O. Yu, Breeder benefits and uranium ore availability, *Nucl. News* **18**, 46–52 (1975).
22. R. Meyer, On the relationship among the utility of assets, the utility of consumption, and investment strategy in an uncertain, but time-invariant world, *in* "OR 69: Proceedings of the Fifth International Conference on Operational Research," (J. Lawrence, Ed.), Tavistock, London (1970).
23. L. L. Pasinetti, Changes in the rate of profit and switches in technique, *Quart. J. Econ.* **80**, 503–517 (1966).
24. L. Pontryagin *et al.*, "The Mathematical Theory of Optimal Processes," Wiley, New York (1962).
25. H. Raiffa, "Decision Analysis: Introductory Lectures on Choices Under Uncertainty," Addison–Wesley, Reading, Mass. (1968).
26. J. Rawls, "A Theory of Justice," Harvard, Cambridge, Mass. (1971).
27. P. Samuelson, A summing up, *Quart. J. Econ.* **80**, 568–583 (1966).
28. P. Samuelson, Some aspects of the pure theory of capital, *Quart. J. Econ.* **51**, 469–496 (1937).
29. J. Satia and R. Lave, Markovian decision processes with uncertain transition probabilities, *Operations Res.* **21**, 728–740 (1973).
30. R. Solow, Intergenerational equity and exhaustible resources, *Rev. Econ. Stud., Symp.* 29–46 (1974).
31. H. Thomas and P. Watermeyer, Mathematical models: A stochastic sequential approach, *in* "Design of Water-Resource Systems," (Maass *et al.*, Eds.), Harvard, Cambridge, Mass. (1962).
32. W. K. Viscusi and R. Zeckhauser, Future mindedness, uncertainty, and irreversibility, Discussion Paper No. 30D, Kennedy School of Government, Harvard Univ. (1975).
33. M. Weinstein and R. Zeckhauser, Use patterns for depletable and recycleable resources, *Review Econ. Studies, Symp.* 67–88 (1974).
34. J. E. Wilen, Irreversibilities and stratospheric flight, Working Paper, Univ. of British Columbia (1974).
35. R. Zeckhauser, Resource allocation with probabilistic individual preferences, *Amer. Econ. Review* **59**, 546–552 (1969).
36. R. Zeckhauser, The risks of growth, *Daedalus* **102**, 103–118 (1973).
37. R. Zeckhauser, M. Spence, and E. Keeler, The optimal control of pollution, Discussion Paper No. 1D, Kennedy School of Government, Harvard Univ. (1970).