Staff Papers Series

P85-26 August 1985

UNCERTAINTY, INFORMATION, AND IRREVERSIBLE INVESTMENTS

by

Theodore Graham-Tomasi

Department of Agricultural and Applied Economics

University of Minnesota
Institute of Agriculture, Forestry and Home Economics
St. Paul, Minnesota 55108

UNCERTAINTY, INFORMATION, AND IRREVERSIBLE INVESTMENTS

by

Theodore Graham-Tomasi

August, 1985

UNCERTAINTY, INFORMATION, AND IRREVERSIBLE INVESTMENTS

by

Theodore Graham-Tomasi[1/]

Introduction

A large literature has developed concerning the evaluation of the social desirability of undertaking actions which irrevocably will alter unique natural environmental resources.

The consideration of uncertainty in this context has led to the concepts of option value and quasi-option value (QOV). The former concept concerns the relationship between ex-ante and ex-post welfare evaluation measures.[2/] The latter concept, with which this paper is concerned, has to do with the impact on investment criteria of opportunities for learning which might eliminate or reduce uncertainty. In particular, Arrow and Fisher (1974) demonstrated that decisions with learning appropriately are more conservative than decisions made when learning possibilities do not exist or are ignored.

While the implications of the analysis are clear, the definition of QOV is not. Arrow and Fisher identified QOV as "a reduction in net benefits from development" (1974, p. 315). Similarly, Henry (1974a) defined an "irreversibility effect," which is a reduction in the expected net benefits of development when uncertainty and learning are recognized relative to the certainty case. However, other authors have identified QOV as the quantity which, if incorporated into the decision rule regarding development used in the non-learning case, would lead to the same decision rule in the learning case (Krutilla and Fisher, 1975, p. 72). This approach is taken by Bernanke (1983) where there are several possible investments. In this latter sense, QOV is seen as a "shadow tax" on development which leads to efficient investment decisions.

Unfortunately, the models employed by Arrow, Krutilla, and Fisher were simple and such that the relationship between these alternative definitions of QOV is obfuscated. In particular, they assumed the the area of land available for development is equal to one and that benefit functions are linear so that either all of the area or none of it is developed.

In a different setting in which the benefit of holding given stocks of developed and preserved land is given by a strictly concave function, and a "zero or one" restriction is not imposed, one is presented by a more obvious choice between definitions of QOV, and the distinction between the "net benefits" and "shadow tax" approaches is clear. In this framework, the former definition is seen to be equivalent to the expected value of information as defined in utility terms by, for example, Gould (1934). Conrad (1980) points out that the two ideas are equivalent. Hanemann (1983) correctly notes that this value is conditional on no development in the first period.

I propose that the Krutilla/Fisher shadow tax definition be given the name quasi-option tax (QOT) and that the term quasi-option value (QOV) be maintained when discussing changes in net benefits. Hence, throughout the paper, I refer to these two concepts without, I hope, risk of confusion.

First, I explore a fairly general model of investment decisions under uncertainty with learning and derive an expression that corresponds to the shadow tax definition offered by Krutilla and Fisher in the linear case.

I show how this differs from the usual definition of the QOV.[3/] In particular, I propose an operational definition of QOT when the choice set is the continuum [0,1], and not equal to {0,1}, as is usual in the literature. Hanemann (1983) notes the distinction between the QOV and shadow tax concepts (defined differently than is done here) in this context and states that the concept, therefore, is not well defined. However, that two distinct concepts exist which are equal in one special case does not obviate their usefulness. Bernanke (1983) proposed a measure similar to that studied here. However, he considers a menu of projects which can be either implemented on deferred and considers the problem of choice from among these alternatives. He incorrectly associates his measure with QOV (this is noted by Hanemann (1983)).

Second, I investigate the properties of the QOV and what I propose as QOT. In particular, I am interested in how QOT responds to an increase in risk as defined by Rothschild and Stiglitz (1970).

The paper is organized as follows. In the next section, I outline the basic model of QOV following Arrow and Fisher (1974). The simplest model is presented and the distinction between QOT and the QOV is drawn. I also show how these can be the same in magnitude in this simple case.

The relationship between the QOV and QOT is obscured in a model with linear benefit functions. In the third section of the paper, I study a model with concave benefit functions and many possible realizations of the state of nature; however, I retain the assumption of two time periods as well as that concerning perfect information. This section serves two purposes. First, it serves as a forum for the demonstration of the intuitively pleasing result that, when benefit functions are strictly concave, less of a natural environmental resource is developed when learning is anticipated than when learning cannot take place. The method of proof is alternative to that presented in Arrow and Fisher (1974). The mode of proof used here provides a clear idea of the difference between the QOV and QOT; this is the second purpose of the section.

The comparison of no information to perfect information is a special case of the comparison of "better" information. In the fourth section of the paper, I examine the effect that an ability to obtain better information has on irreversible investment decisions. In other respects, the model is the same as that studied in the previous section. The approach taken here is to apply the concept to a "more informative" experiment as found in Marschak and Miyasawa (1968) in a dynamic setting via dynamic programming arguments. The exposition here follows closely that of Epstein (1980).

The fifth section extends the inquiries of the previous section to a multi-period model. The extension is straightforward and requires only a change of interpretation. The final section is a discussion.

Quasi-Option Tax and Quasi-Option Value

The existence of QOT and its relationship to quasi-option value (QOV) in the linear case easily is demonstrated in a model with two periods, two possible states of nature, and perfect information. The presentation here follows Arrow and Fisher (1974) conceptually, but the details of their analysis are modified to provide a closer tie to subsequent investigations in this chapter.

There is a tract of homogeneous underdeveloped area the size of which is normalized to one. The amount of new development in period t is q_t for $t = 1,2$. The net benefits of development and preservation per unit of area per unit of time are D_t and P_t, respectively. It is assumed that, while D_1 and D_2 are known, the values of P_t are not known. As with models of investment by a competitive firm under uncertainty, the timing of any resolution of uncertainty relative to the time when input choices must be made is crucial in the determination of the results of the analysis.[4/] Here, I assume that P_1 is known when q_1 must be chosen. I then compare two cases, one in which no information concerning P_2 is forthcoming, and one in which P_2 is known with certainty, before q_2 must be chosen.

I assume that *ex ante* the decision-maker (DM) believes (correctly) that P_2 is "small" with probability Π and that it is "large" with probability $1-\Pi$.[5/] The realization P_2^s is small in that $D_2-P_2^s$ is strictly positive; P_2^h is large in that $D_2-P_2^h$ is strictly negative.

First I examine the case in which no new information about P_2 is available to the DM until after q_2 must be decided upon. Then both q_1 and q_2 are chosen at the initial date. The DM is assumed to be risk-neutral; thus, the objective function is the expected present value of net benefits. The DM's point estimate of P_2, conditioned on information available at the beginning of the program, is $\bar{P}_2$, i.e., $\bar{P}_2 = \Pi P_2^s + (1-\Pi)P_2^h$. The DM solves

$$\max_{q_1,q_2} \quad D_1q_1 + (1-q_1)P_1 + a[D_2(q_1+q_2) + (1-q_1-q_2)\bar{P}_2]$$

$$\text{s.t.} \quad q_1\varepsilon[0,1], \ q_2\varepsilon[0,1-q]$$

where $a=(1+i)^{-1}$ and i is the decision-maker's rate of time preference.

Define $V_1=D_1-P_1$ and $\bar{V}_2=D_2-\bar{P}_2$. Clearly, the optimal plan includes for q_2

$$q_2^* = \begin{cases} 0 \text{ if } \bar{V}_2 < 0 \\ 1-q_1 \text{ if } \bar{V}_2 > 0 \end{cases} \quad \text{all } q_1 \varepsilon [0,1], \tag{1}$$

Turning to the decision in the first period, the decision rule is given by

$$q_1^* = \begin{cases} 1 \text{ if } \bar{V}_2 < 0 & V_1 > |a\bar{V}_2| \\ 0 \text{ if } \bar{V}_2 < 0 & V_1 < |a\bar{V}_2| \\ 1 \text{ if } \bar{V}_2 > 0 & V_2 > 0 \\ 0 \text{ if } \bar{V}_2 > 0 & V_1 < 0. \end{cases} \tag{2}$$

Now, suppose that the decision-maker recognizes that the true value of P_2 will be revealed at the beginning of period two. In this case, the decision on q_2 can be made after P_2 is known. Let V_2 denote the true value of D_2-P_2. The decision rule on q_2^s (the superscript on the variable denotes that it is chosen via a sequential decision procedure) is given by (1) replacing $\bar{V}_2$ by V_2. The decision rule concerning q_1^s is derived by comparing the expected payoff with $q_1^s > 0$ to the expected payoff with $q_1^s = 0$ when the optimal decision rule for q_2^s is used. In the former case the expected payoff is

$$q_1^sD_1 + (1-q_1^s)P_1 + a\ \{\Pi D_2 + (1-\Pi)[q_1^sD_2 + (1-q_1^s)P_2^h]\}. \tag{3}$$

If, on the other hand, $q_1^s=0$, the expected payoff is

$$P_1 + a\ [\Pi D_2 + (1-\Pi)P_2^h]. \tag{4}$$

Subtracting (4) from (3) yields

$$q_1^sD_1 - q_1^sP_1 + a(1-\Pi)q_1^sD_2 - a(1-\Pi)q_1^sP_2^h. \tag{5}$$

Clearly, $q_1^{s*} = 1$ if (5) is positive, while $q_1^{s*} = 0$ if (5) is negative.

Rearranging, the rule is

$$q_1^{s*} = \begin{cases} 1 \text{ if } V_1 > a(1-\Pi)(P_2^h - D_2) \\ 0 \text{ if } V_1 < a(1-\Pi)(P_2^h - D_2). \end{cases} \tag{6}$$

By assumption, $a(1-\Pi)(P_2^h-D_2)$ is a positive number. The existence of QOV is demonstrated by comparing (2) and (6); i.e., by comparing decision rules when opportunities for learning are ignored with decision rules for sequential decisions. Suppose $\bar{V}_2 > 0$. Then from the last two lines of the RHS of (2) we have

$$q_1^* = 1 \text{ if } V_1 > 0$$
$$q_1^* = 0 \text{ if } V_1 < 0$$

as the decision rule ignoring learning, and (6) in the sequential case. When opportunities for learning exist, V_1 must be higher for development to be indicated than it must be if these opportunities either do not exist or are ignored. The amount by which initial development benefits must be higher is just $a(1-\Pi)(P_2^h-D_2)$, the expected present value of the loss if an irreversible decision is undertaken and the state of nature that obtains is such that the decision-maker would reverse the decision if (s)he could; this number is the QOT.

It is apparent that this "shadow tax" approach to QOT bears a resemblance to the concept of the value of information. Here, I follow Gould (1974) in defining the expected value of perfect information. A decision-maker chooses a decision $x \varepsilon X$ and obtains a payoff of $U(x,z)$ where $z \varepsilon Z$ is a random variable with distribution function $F(z)$. In the absence of information, the decision-maker maximizes the expected payoff, i.e., (s)he solves

$$\max_{x \varepsilon X} \int_Z U(x,z)dF(z). \tag{7}$$

Let $\hat{x}$ denote the solution to this problem. Now, suppose that the decision-maker can obtain perfect information about z before choosing x. Let x_z^* denote the optimal choice in this case. Then, before the information actually is obtained the expected payoff with information is

$$\int_Z U(x_z^*, z)dF(z). \tag{8}$$

The expected value of perfect information is defined as the difference between (8) and (7) evaluated at $\hat{x}$, i.e.,

$$VOI = \int_Z U(x_z^*, z)dF(z) - \int_Z U(\hat{x}, z)dF(z). \tag{9}$$

It is easy to show that VOI $\geq$ 0 (see Gould, 1974, p. 67). This makes intuitive sense, as the result says that obtaining information costlessly can never make the decision-maker worse off.

To compare this definition to the concept of QOT developed above, I introduce the following notation. Let $E_oPV(1)$ be the expected present value of the payoff in the problem discussed previously when opportunities for learning are ignored and the decision in the first period is to develop all of the area. Similarly, define $E_oPV(0)$ as this present value when the initial decision is to develop none of the area. When a sequential decision procedure is used, denote the expected payoffs by $E_sPV(1)$ and $E_sPV(0)$.

Suppose that $E_oPV(1) > E_oPV(0)$ and the "ignorance case" yields full development of the area at the initial date. The lesson of the previous analysis is that even if this is the case, it may be that $E_sPV(0) > E_sPV(1)$, i.e., even if it is optimal to develop in the ignorance case it may be optimal to refrain from development in the learning case. Assume that these two inequalities hold.

Since the definition of VOI forwarded above uses the expected payoffs evaluated at optimal choices, with the assumption of the previous paragraph, one would define

$$VOI = E_s PV(0) - E_o PV(1). \tag{10}$$

Applying this definition to the problem studies in the previous section,

$$\begin{aligned} VOI &= P_1 + a[\Pi D_2 - (1-\Pi)P_2^h] - D_1 - aD_2 \\ &= P_1 - D_1 + a(1-\Pi)(P_2^h - D_2). \end{aligned} \tag{11}$$

Recalling the discussion above, the second term on the RHS of (11) is the QOT for the problem and it is immediate that, in general, QOT≠VOI.

One difference in the two concepts arises from the fact that the QOV allows the choice variable to take on its maximizing value in both the learning and ignorance cases. On the other hand, QOT in this linear case seems to be the difference between the expected payoff with learning and the expected payoff under ignorance when $q_1=0$ in both cases.

In this linear model, QOT restricts the assessment to the choice $q_1=0$ in both cases, i.e.,

$$QOT = E_s PV(0) - E_s PV(0). \tag{12}$$

Subtracting (10) from (12),

$$QOT - VOI = E_o PV(1) - E_o PV(0),$$

or, more suggestively,

$$QOT - VOI = E_o PV(q_1^*) - E_o PV(0). \tag{13}$$

In the linear case, equality will hold in (13) if the optimal choices in the two cases coincide. Implicitly, Conrad assumed $q_1^*=q_1^{s*}=0$ and this explains why he was able to present a numerical example with QOT=VOI even though the two concepts are distinct.

The two ideas are associated with different questions. The VOI examines the difference in expected total payoff in the learning and ignorance cases and hence is concerned with the expected value of the indirect utility function. QOT, on the other hand, examines the bias that arises in the decision rule used to make initial development choices, i.e., QOT is a marginal concept. It seeks the adjustment to initial development benefits in the learning case that would lead to the same level of initial development as in the ignorance case.

One problem here is that, in a linear model, totals and marginals easily are confused when the upper bound on initial development is chosen to be one. The difference between QOT and VOI is more readily seen in a model with strictly concave benefit functions. In the next section of the paper I study such a model, while retaining the two-period and perfect information assumptions.

A More General Two-Period Model

In this section I change my notation somewhat. As before, q_t is the amount of new development in period t. Now, define $D_t = \Sigma_t q_t$ and $P_t=1-D_t$; here, D and P are areas of land and not unit benefits as in the previous section. The net benefits of any given level of development are given by the function $U(D_t,P_t,z_t)$ where z_t is a parameter that cannot be observed by the DM. I assume that, for any given value of z, U is thrice continuously differentiable, strictly increasing in D and in P, and strictly concave in (D,P) on all of its domain, the closed unit interval. Moreover, for convenience, I assume

$$\lim_{D\to 0} \partial U/\partial D = \infty = \lim_{P\to 0} \partial U/\partial P \tag{14}$$

Perfect information is obtained by performing an experiment Y which provides a message y that is perfectly correlated with z; that is, observation of y tells the DM exactly what z*, the true value of z, is. Obviously, this is a special case of an experiment which provides noisy information about z, the implications of which will be addressed in the following section of the chapter.

As in the previous section, z* is known to lie in a set Z. The DM's subjective beliefs about z are summarized by the prior probability distribution function F(z). Here, the set Z could contain either a finite, countable, or uncountably infinite number of elements.

Consider first the case with no experimentation in which the message system Y either doesn't exist, or is ignored. The DM solves

$$\max_{q_1,q_2} \sum_{t=1}^{t=2} a^{t-1} \int_Z U(D_t,1-D_t,z)dF(z)$$

$$\text{s.t.} \quad D_t = \Sigma q_t$$

$$D_t \varepsilon [0,1]$$

$$q_t \geq 0.$$

Since the problem is stationary, the DM chooses q to solve

$$\max \int_Z U(q,1-q,z)dF(z).$$

By assumption, U is linear in z; hence, by Jensen's Inequality, this problem is equivalent to

$$\max \quad U(q,1-q, \int_Z zdF(z)) = U(q,1-q,\bar{z}). \tag{16}$$

Let $\bar{q}$ be a solution to this problem; previous assumptions ensure that $\bar{q}$ exists and is unique. The optimal policy for the non-experimenting DM, then, is to set $(q_1,q_2)=(\bar{q},0)$.

Next, consider the problem in which experimentation reveals $z*$ before q_2 is chosen. In period 2 the DM solves

$$\max_{0 \leq q_2 \leq 1-q_1} \quad U(q_1+q_2,1-q_1-q_2z*). \tag{17}$$

Denote the solution (which exists and is unique) to this problem by $q_2^*(q_1,z*)$, and define the indirect utility function

$$V(q_1,z*) = U(q_1+q_2^*(q_1,z*),1-q_1-q_2^*(q_1,z*),z*). \tag{18}$$

Lemma 1. $V(q,z)$ is differentiable and non-increasing in q.

Proof: That V is differentiable follows from the Implicit Function Theorem, which implies that $q*(\cdot)$ is a $C^{(1)}$ function. To show that V is non-increasing, note that, by (14) $q_2^* < 1-q_1$. Hence, the conditions (necessary and sufficient) characterizing q_2^* are

$$\frac{\partial U}{\partial q_2} \; (q_1 + q_2^*, \; 1-q_1 - q_2^*,z*) \leq 0 \tag{19a}$$

$$q_2^* \; [\frac{\partial U}{\partial q_2} \; (q_1 + q_2^*, \; 1-q_1 - q_2^*,z*)] = 0. \tag{19b}$$

q.e.d.

Lemma 1 is quite intuitive. If, having inherited q_1, more development in the second period is undertaken, then a marginal change in q_1 does not affect the DM's utility, since q_1 and q_2 are perfect substitutes.

If, on the other hand, q_2^* is zero given q_1, the non-negativity constraint is binding. Then, increasing q_1 tightens this constraint making the DM worse off.

Given that an optimal policy is to be used in the second period, standard dynamic programming arguments (i.e., the Principle of Optimality) tells us that the solution to the first-period investment problem is obtained by solving

$$\max_{q_1} \int_Z U(q_1,1-q_1,z)dF(z) + a \int_Z V(q_1,z)dF(z) \qquad (20)$$

Denote the solution to this problem by q_1^*. The key result of this section is the demonstration that, if learning is anticipated, less development is undertaken in the first period.

Theorem 1. $q_1^* \leq \bar{q}$. Further, if there exists $\hat{z}$ such that $U_D(\bar{q},1-\bar{q},\hat{z}) - U_P(\bar{q},1-\bar{q},\hat{z}) < 0$ and $F'(\hat{z}) > 0$, then $q_1^* < \bar{q}$.

Proof: Suppose the contrary, i.e., that $q_1^* > \bar{q}$. Since $\bar{q}$ is the unique maximizer of first period expected payoffs,

$$\int_Z U(\bar{q},1-\bar{q},z)dF(z) > \int_Z U(q_1^*,1-q_1^*,z)dF(z).$$

By Lemma 1, $V(\bar{q},z) \geq V(q_1^*,z)$ for all z. Hence,

$$\int_Z V(\bar{q},z)dF(z) \geq \int_Z V(q_1^*,z)dF(z).$$

Adding these provides

$$U(\bar{q},1-\bar{q},\bar{z})+a \int_Z F(\bar{q},z)dF(z) > \int_Z U(q_1^*,1-q_1^*,z)dF(z)$$

$$+a \int_Z V(q_1^*,z)dF(z).$$

But, by assumption, q_1^* is maximal, a contradiction.

To prove the second part, the first order conditions for the problem (20) are, since q_1^* is interior to [0,1] by (14),

$$U_D(q_1^*,1-q_1^*,\bar{z})-U_P(q_1^*,1-q_1^*,\bar{z}) + a \int_Z V_{q_1}(q_1^*,z)dF(z) = 0. \tag{21}$$

By hypothesis the last term on the LHS is strictly negative since $V_{q_1} \leq 0$ for all z and $V_{q_1} < 0$ for $\hat{z}$, where z arises with positive probability. Recalling the definition of $\bar{q}$,

$$U_D(\bar{q},1-\bar{q},\bar{z}) - U_P(\bar{q},1-\bar{q},\bar{z}) = 0.$$

The result follows from the strict concavity of U.

q.e.d.

Now, consider the Krutilla/Fisher shadow-tax definition of QOT. Suppose a DM who ignores opportunities for learning when they exist and who acts myopically to maximize first period expected payoffs must pay a tax of τ per unit of development undertaken. This DM solves

$$\max_q \; U(q,1-q,\bar{z}) - \tau q.$$

The solution to this problem is, of course,

$$U_D(q,1-q,\bar{z}) - U_P(q,1-q,\bar{z}) - \tau = 0. \tag{22}$$

Clearly, if the tax is set such that

$$\tau = -a \int_Z \frac{\partial V}{\partial q_1}(q_1,z)dF(z),$$

the DM who ignores learning is led to an efficient investment decision since, in this case, (22) and the first order conditions for (20) (given by

(21)) coincide. Since U(•) is strictly concave, the solutions to the problems must coincide as well. Thus, since I propose that QOT be defined as in Krutilla and Fisher (1975), I propose

$$QOT \equiv -a \int_Z \frac{\partial V}{\partial q_1}(q_1,z)dF(z).$$

It provides some intuition into this definition if it is noted that, while the integral is taken over all of Z, the derivative under the integral is zero for a subset of Z and the numerical value of QOT is unchanged if this set is deleted. Define

$$\bar{Z} = \{z: \; U_D(\bar{q},1-\bar{q},z)-U_P(\bar{q},1-\bar{q},z) < 0\}.$$

The set $\bar{Z}$ is the subset of Z such that, if $z^* \varepsilon \bar{Z}$ and the myopically profitable level of development is undertaken in the first period, then no new development would be undertaken in the second period. Hence,

$$-QOT = a \int_{\bar{Z}} \frac{\partial V}{\partial q_1}(q_1,z)dF(z).$$

This formulation makes it clear that QOT is the expected present value of the second period loss if an irreversible decision is implemented at its myopically profitable level, where the loss is averaged over the possible states of nature under which the DM would reverse his/her decision if he/she could. This, of course, is just the interpretation given in the previous section of the paper. It is apparent that the two formulations are equivalent.

This also makes clear the role played by the assumption which leads to the conclusion that $q_1^* < \bar{q}$, i.e., that $\bar{Z}$ has a non-empty subset with positive measure. Obviously, if it is not possible to find out that an <u>ex ante</u> decision is incorrect, then the prospect of "learning" will not affect that decision. Equally obviously, the VOI in this case is zero as well.

Further intuition can be gained with the aid of a diagram. In Figure 1 the marginal benefit functions are depicted. The myopic investment level q^m occurs where expected marginal net benefits of investment are zero, i.e., where $\partial U/\partial D = \partial U/\partial P$. If, in the learning case, there exists a value of z such that reversing some initial development would be desired, and this z arises with positive probability, then QOT is positive. QOT drives a wedge between marginal first period development and preservation benefits, and it is optimal to develop less of the area in the first period, as shown.

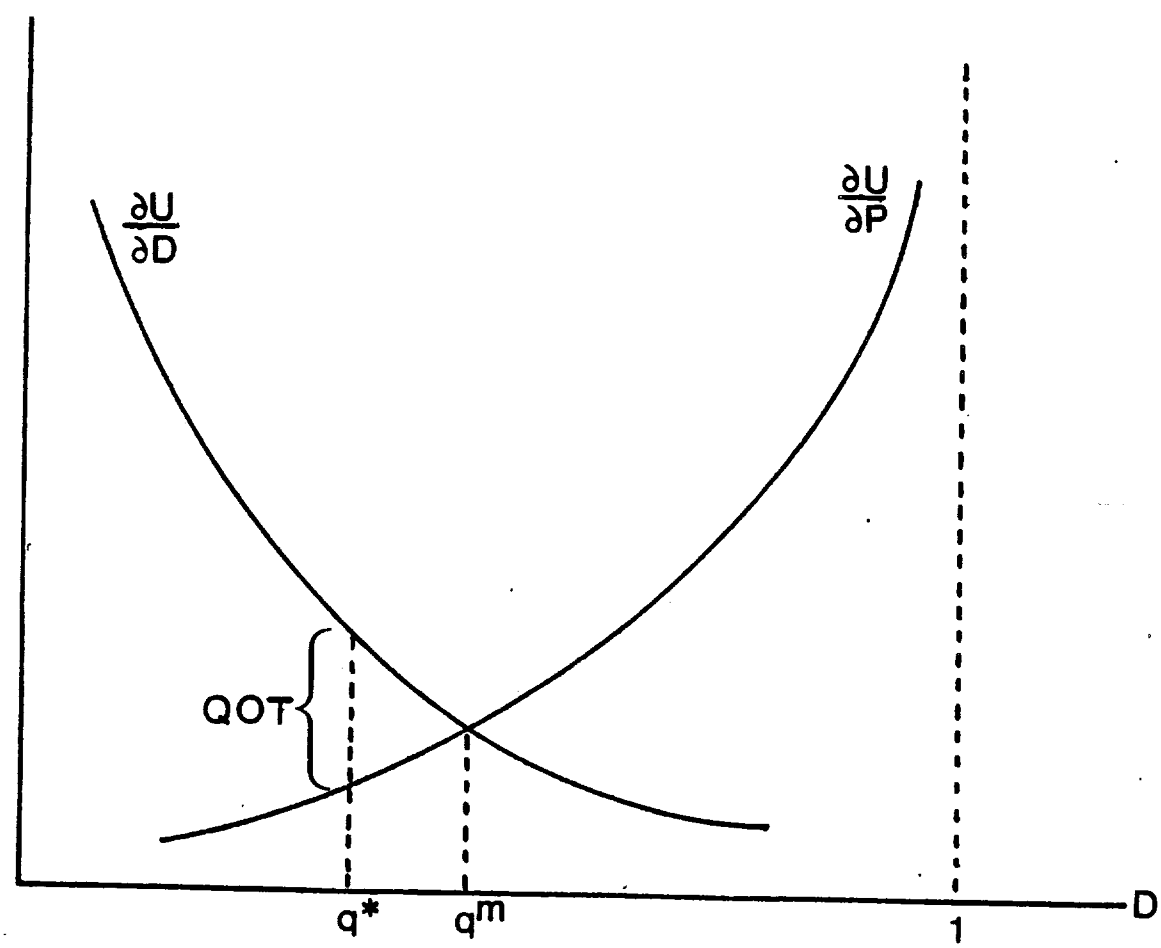

FIGURE 1. Optimal and Myopic Investment

In the model outlined above, I assumed that the utility function and unknown parameter of both are stationary. In the context of the deterministic natural environments literature discussed in the introduction of the paper, this assumption may not be warranted. To incorporate the possibility of preferences changing over time, I now allow z to vary through time. The most important change in the analysis is that, in general, the optimal policy in the non-learning case no longer is of the form $(\bar{q},0)$.

Let the probability distribution on z_t be denoted by $F^t(z_t)$ and define

$$\bar{z}_t = \int_Z z_t F^t(z_t).$$

The non-experimenting DM solves

$$\max_{q_1,q_2} \; U(q_1,1-q_1,\bar{z}_1) + aU(q_1+q_2,1-q_1-q_2,\bar{z}_2). \tag{23}$$

Let $(\bar{q}_1,\bar{q}_2)$ be the solution to (23) and let $\bar{J}(\bar{q}_1,\bar{q}_2)$ be the indirect objective function. Naturally, the solution vector satisfies (recall (14))

$$\frac{\partial\bar{J}}{\partial q_1} = 0 \tag{24}$$

$$\frac{\partial\bar{J}}{\partial q_2} \leq 0 \qquad \bar{q}_2[\frac{\partial J}{\partial q_2}] = 0 \tag{25}$$

It is important to note that the problem (23) is a two-period version of the problem considered by Fisher, Krutilla, and Cicchetti (1972). Suppose that $\bar{q}_2 > 0$. Then, $\partial\bar{J}/\partial q_2 = 0$. Using (24) and the definition of $\bar{J}$, it is immediate that

$$U_D(\bar{q}_1,1-\bar{q}_1,\bar{z}_1) - U_P(\bar{q}_1,1-\bar{q}_1,\bar{z}_1) = 0 \tag{26}$$

But this is just the condition satisfied by the myopic level of initial development q^m where marginal net benefits of first period development are zero.

This discussion indicates that $\bar{q}_1=q^m$ if, given q^m, it is expected that more development would be indicated in the second period of the program. One might call this an "expected free interval" since the DM is free to invest at the myopically profitable level in the first period. If, on the other hand, one expects $\bar{q}_2=0$ given $q_1=q^m$, then

$$U_D(\bar{p}_1,1-\bar{q}_1,\bar{z}_1)-U_P(\bar{q}_1,1-\bar{q}_1,\bar{z}_1) + \partial\bar{J}/\partial q_2 = 0$$

and it follows that $\bar{q} \leq q^m$. Disinvestment ($q_2 < 0$) is desired, but this action is blocked by the irreversibility constraint. Along this "expected blocked interval" current investment is curtailed to a level smaller than the myopic level of investment. This is the Arrow (1968) result that it may be optimal to refrain from myopically profitable investment if disinvestment, which is impossible, would be desired in the near future.

Naturally, if the DM is risk neutral and no learning is anticipated, it is not surprising that the DM may replace the unknown parameter z by its expected value in each period and proceed to undertake a Fisher, Krutilla, Cicchetti benefit-cost analysis. The crux of this section of the chapter is to show that this certainty equivalence result does not hold when learning is anticipated. I wish to show that $q_1^* \leq \bar{q}_1$.

There are three possible cases to consider. The first case is trivial: Suppose that, for all z_2 that arise with positive probability, one has

$$U_D(q_1^m,1-q_1^m,z_2) - U_P(q_1^m,1-q_1^m,z_2) > 0 \tag{27}$$

where q_1^m is the myopic investment level for the first period, i.e. q_i^m is defined by

$$U_D(q_i^m, 1-q_i^m, \bar{z}_i) - U_P(q_i^m, 1-q_i^m, \bar{z}_i) = 0 \quad i=1,2 \qquad (28)$$

Then, $q_2^* = q_2^m - q_1$ for all z_2 and clearly $q_1^* = q_1^m$. This is the case where there is nothing to be learned and $q_1^* = \bar{q}_1 = q_1^m$. One might call this a "sure free interval."

The second case is equally as trivial as the first; here, assume $q_1^m < q_2^m$. Then $\bar{q}_1 = q_1^m$ and Theorem 1 holds, since subscripting z by t appropriately does not affect its proof. This is the expected free interval case.

The third case is more interesting; further, it corresponds to the situation envisioned by Fisher, Krutilla, and Cicchetti (1972). In this case an expected blocked interval exists, i.e., $q_1^m > q_2^m$. Now, by the discussion above, $\bar{q}_1 < q_1^m$. The proof of Theorem 1 no longer applies since it cannot be concluded that $\bar{q}_1$ maximizes first period utility.[6/]

I will defer the proof of the proposition that $q_1^* \leq \bar{q}_1$ in this case until the next section of the paper. There I show that $q_1^* \leq \bar{q}_1$ in a more general case of less than complete resolution of uncertainty when learning takes place. The third case here naturally holds as a special case of that more general model.[7/]

Imperfect Information

To characterize the uncertainty in the problem and the possible receipt of information, I follow the approach of Marschak and Miyasawa (1968).[8/] In particular, I assume that the decision-maker holds subjective beliefs about the true value of the parameter z_t, beliefs given by the prior probabilities $(\Pi_{t-1}, \ldots, \Pi_{tn})$, where

$$\Pi_{ti} = \text{Prob}(z_t = z_{ti})$$

The decision-maker is able to perform an experiment Y, which provides information about the true value of z_2. The set of possible messages corresponding to Y is $(y_1, \ldots, y_m)$. The tie between the experiment and the parameter is provided by the likelihood matrix

$$\Delta = [\delta_{ij}] = \text{Prob}(Y=y_i \mid z_2=z_{2j}).$$

Then, the posterior probabilities are found via Bayes' Theorem. The matrix of posterior probabilities is denoted by

$$\theta = [\theta_{ij}] = \text{Prob}(z_2=z_{2i} \mid Y=y_j).$$

Let the joint distribution of Y and z_2 be β_{ij} and the marginal distribution of Y be λ_j, i.e., $\lambda_j=\text{Prob}(Y=y_j)$. Then

$$\lambda_j = \Sigma_i \beta_{ij} = \Sigma_i \Pi_{2i} \delta_{ij}$$

and

$$\beta_{ij} = \delta_{ij} \Pi_{2i} = \theta_{ij} \lambda_j.$$

If a particular message y_j is obtained from performing the experiment, then the conditional distribution of z_2 is provided by the j^{th} column of the matrix $[\theta_{ij}]$; this can be obtained via knowledge of either the joint distribution or the likelihood matrix.

Consider a general decision problem in which a payoff of $w(x,z)$ is realized if a decision $x \varepsilon X$ is chosen and the true state of the world is z. If a message of y is received, the decision-maker chooses $x^*(y)$, where

$$\Sigma_j \Sigma_i w(x^*(y),z) \beta_{ji} \geq \Sigma_j \Sigma_i w(x,z) \beta_{ji} \text{ for all } x \varepsilon X,$$

i.e., the decision-maker chooses x* to maximize the expected payoff. Define

$$U^*(\beta) = \Sigma_j \Sigma_i w(w^*(y),x) \beta_{ji}.$$

By the discussion above,

$$U^*(\beta) = \Sigma_j \lambda_j \max_x \Sigma_i \theta_{ij} w(x,z).$$

Marschak and Miyasawa (1968) compare the information content of two alternative experiments Y and Y' and offer the following basic definition: Y is more informative than Y' if $U^*(\beta) \geq U^*(\beta')$. It is worthwhile noting that comparing "no information" to "perfect information" is a special case of the general "more informative" comparison defined above. If the definition of the value of information given in the second section is extended to the imperfect information setting, one has

$$VOI(Y) = U^*(\beta) - U^*(\beta^o)$$
$$VOI(Y') = U^*(\beta') - U^*(\beta^o)$$

where β^o is such that $\theta^o_{ij}=\Pi_{2i}$, that is, the message and the value of the parameter are independent. The experiment Y is defined as more informative than Y' if its information content has higher value.

Returning to the development versus preservation problem at hand, the decision-maker solves

$$\max_{q_1,q_2} E\{U_1(D_1,P_1,z_1) + aU_2(D_2,P_2,z_2)\}$$

$$\text{s.t. } q_1 \in (0,1),\ q_2 \in [0,1-q_1).$$

Consider the development problem in the second period. Given a value of q_1 and having observed a message y_j, the problem is

$$\max_{q_2 \varepsilon [0,1-q_1]} \Sigma_i \; \theta_{ij} \; U_2(q_1+q_2, 1-q_1-q_2, z_{2i})$$

Since I have assumed that the marginal utility of wilderness gets infinitely large as the stock of wilderness goes to zero, the solution to this problem is strictly less than $1-q_1$; however, a zero solution is admitted. Define the Lagrangean function

$$L(q_2,\gamma;q_1,\theta_j) = \Sigma_i \; \theta_{ij} \; U_2(q_1+q_2,1-q_1-q_2,z_{2i}) + \gamma \; q_2. \qquad (27)$$

Since the problem is one of concave programming and U is strictly concave, the first order necessary conditions are sufficient for a unique global maximum. These conditions are

$$\Sigma_i \theta_{ij} \; [\frac{\partial U_2}{\partial D} (q_1+q_2^*,1-q_1-q_2^*,z_{2i}) - \frac{\partial U_2}{\partial P} (q_1+q_2^*,1-q_1-q_2^*,z_{2i})] + \gamma^* = 0 \qquad (28)$$

$$\gamma^* \geq 0, \; q_2^* \geq 0, \; \gamma^* q_2^* = 0 \qquad (29)$$

Denote the solution to this second-period problem by $q_2^*(q_1,\theta_j)$, i.e., the function q_2^* solves (28) identically. Further, let

$$L^*(q_1,\theta_j) = L(q_1,q_2^*(q_1,\theta_j),\theta_j).$$

By the complementary slackness condition (19b), L* is identically equal to the expected value of the indirect utility function.

The development decision in the first period is obtained by solving

$$\max_{q_1 \varepsilon (0,1)} \Sigma_i \; \Pi_{1i} \; U_1(q_1,1-q_1,z_{1i}) + a \; \Sigma_j \; \lambda_j \; L^*(q_1,\theta_j).$$

The concavity of L* in q_1 follows from the concavity of U_2; hence, since U_1 is strictly concave, the first order necessary conditions are sufficient for a unique global maximum. These are

$$\Sigma_i \Pi_{1i} [\frac{\partial U_1}{\partial D} (q_1, 1-q_1, z_{1i}) - \frac{\partial U_1}{\partial P} (q_1, 1-q_1, z_{1i})] + a \Sigma_j \lambda_j \frac{\partial L^*}{\partial q_1} = 0. \tag{30}$$

Let $q_1^* (\Pi, \lambda)$ denote the solution to this equation. The key to the analysis of q_1^* is the last term in (30). Note that, by the envelope theorem,

$$- \frac{\partial L^*}{\partial q_1} = \gamma \geq 0.$$

Hence,

$$\frac{\partial L^*}{\partial q_1} (q_1, \theta_j) \leq 0.$$

Note that $\gamma = 0$ if $q_2^*(q_1, \theta_j) > 0$. That is, the shadow price of passing a marginal unit of undeveloped land to the next period is zero if that land would be developed in that period.

Let q_1^m be defined by

$$\Sigma_i \Pi_{1i} [\frac{\partial U_1}{\partial D} (q_1^m, 1-q_1^m, z_{1i}) - \frac{\partial U_1}{\partial P} (q_1^m, 1-q^m, z_{1i})] = 0. \tag{31}$$

If the decision-maker solves a static problem in the first period so as to maximize the expected value of utility while ignoring irreversibility and opportunities for learning, the optimal myopic level of development is q_1^m.

So far, the analysis is exactly as in the previous section where z is non-stationary except that here I have restricted the analysis to a finite (or countable) number of possible states of nature.

I next inquire into the effect that the opportunity to perform a "more informative experiment" has on optimal first period development decisions, where more informative is defined as in Marschak and Miyasawa (1968). The key result is the following

Lemma 2. If Y is more informative than Y', then

$$\Sigma_{j=1}^{m} \lambda_j \ \rho(\theta_j) \geq \Sigma_{j=1}^{m} \lambda'_j \ \rho(\theta'_j)$$

for any convex function ρ defined on the set of vectors forming an (m-1) dimensional simplex.

Proof: Marschak and Miyasawa (1968), Theorem 12.1.

Thus, if it can be demonstrated that $\partial L^*/\partial q_1$ is concave in θ_j, the result $q_1^* \leq \bar{q}_1$ would follow immediately from Lemma 2. Before proceeding, it is worthwhile pointing out that the number $\Sigma_j \lambda_j \ \partial L^*/\partial q_1$ is the QOT for the model at hand. Thus, Lemma 2 would indicate that QOT is an increasing function of the informativeness of the experiment to be performed.

It is worthwhile pointing out that the value function L* itself is convex in θ_j.

Theorem 2. $L^*(q_1,\theta_j)$ is convex in θ_j.

Proof: Define

$$\Psi(q_2;q_1,\theta_j) = \Sigma_i \ \theta_{ij} \ U(q_1+q_2,1-q_1-q_2,z_{2i}),$$

and let q_2^* solve

$$\max_{q_2} \ \Psi(q_2;q_1,\theta_j).$$

Let $x\varepsilon[0,1]$ and define $\theta\alpha = \alpha\theta_j + (1-\alpha)\theta_k$ for $j \neq k$. Clearly,

$$\Psi(q_1;q_2,\theta\alpha) = \alpha\Psi(q_2;q_1,\theta_j) + (1-\alpha)\Psi(q_2;q_1,\theta_k).$$

Now, in general,

$$\sup \{f+g\} \leq \sup \{f\} + \sup \{g\}.$$

Hence,

$$\sup \{\Psi(q_2;q_1,\theta^{\alpha}\} = \sup \{\alpha\Psi(q_2;q_1,\theta_j) + (1-\alpha)\Psi(q_2;q_1,\theta_k)\}$$

$$\leq \alpha \sup \{\Psi(q_2;q_1,\theta_j)\} + (1-\alpha) \sup \{\Psi(q_2;q_1,\theta_k)\}$$

$$= \alpha \Psi(q_2^*;q_1,\theta_j) + (1-\alpha)\Psi(q_2^*;q_1,\theta_k),$$

whence

$$\Psi(q_2^*;q_1,\alpha\theta_j + (1-\alpha)\theta_k) \leq \alpha\Psi(q_2^*;q_1,\theta_j) + (1-\alpha)\Psi(q_2^*;q_1,\theta_k).$$

q.e.d

The implication of this is that one could proceed backwards, using Lemma 2 as a definition, and establish that more informative experiments have an information content that has a higher value as Gould (1974) did for the special case of perfect information.

The proof that $\partial L^*/\partial q_1$ is concave brings out some interesting implications of the model of QOT.

Lemma 3. $\partial L^*/\partial q_1$ is concave in θ_j.

Proof: The envelope theorem informs us that for small perturbations of q_1, the total change in the indirect objective function can be calculated without accounting for changes through the choice variable, i.e. (ignoring the time subscript on U),

$$\partial L^*/\partial q_1 = \Sigma_i \; \theta_{ij}[U_D(q_1+q_2^*,1-q_1-q_2^*,z_{2i}) - U_P(q_1+q_2^*,1-q_1-q_2^*,z_{2i})].$$

From the first order conditions (28) and complementary slackness, if $q_2^* > 0$, then $\partial L^*/\partial q_1 = 0$. However, if $q_2^* = 0$, then

$$\partial L^*/\partial q_1 = \Sigma_i \; \theta_{ij}[U_D(q_1,1-q_1,z_{2i}) - U_P(q_1,1-q_1,z_{2i})] \leq 0$$

This function is linear in θ. Thus, the function $\partial L^*/\partial q_1$ is concave and piecewise linear.

q.e.d.

This lemma is the key step is establishing the basic result of this section.

Theorem 3. If a more informative experiment is to be performed, then less of the resource is developed in the first period.

Proof: Since $\partial L^*/\partial q_1$ is concave by Lemma 3, QOT is convex by definition. Thus, by Lemma 2, a more informative experiment raises the QOT. The optimal first period investment decision is given by

$$U_D - U_P = QOT.$$

Since U is assumed to be strictly concave in (D,P), the result follows.

q.e.d.

This result is congenial to intuition, for it says that as more uncertainty can be reduced in the future, flexibility becomes more valuable. This is an obvious generalization of the perfect information result established in the previous section of the chapter.

It is important to note that if two message systems are compared and both of these systems give rise to expected values of z_2 such that the non-negativity constraint is either binding or not binding, then no change in first period behavior is indicated. This follows since both potential messages lie on the same linear segment of $\partial L^*/\partial q_1$. By Jensen's Inequality, such a segment is both concave and convex and, from Lemma 2, $\partial L^*/\partial q_1$, and hence, q_1^*, is unchanged. If, on the other hand, one message corresponds to a binding constraint in the second period and one corresponds to a slack constraint, then strictly less development in the first period is optimal.

It is of interest to inquire into how both the VOI and QOT respond to an increase in prior uncertainty regarding the parameter z, where an increase in uncertainty is defined as a mean preserving spread (MPS) of the distribution function of z. Formally, an MPS is defined as:

Definition: Let X, Y, and Z be real random variables such that $E(Y|Z) = 0$. The X is an MPS of Z if X has the same distribution as the joint distribution of Y+Z.

The key result of the arguments used here has been provided by Rothschild and Stiglitz (1970):

Lemma 4. Let f(b) be a real valued function. If f is convex (concave) then the expected value of f is not decreased (not increased) when its argument is subjected to an MPS, with strict inequality holding if the curvature is strict.

Gould (1974) has established that if U is linear in z, the VOI increases when F(z) is subjected to an MPS. Gould's proof is straightforward. Recall the definition of the VOI given in (9):

$$VOI = E[U(X_z^*,z)] - E[U(\hat{x},z)].$$

The last term on the RHS of (9) is linear and, therefore, by Lemma 4, invariant to an MPS of z. Thus, if the first term on the RHS of (9) is convex in z, an MPS of z increases its value and, thereby, increases the VOI as well. If this term is differentiated twice with respect to z, one finds that it is, indeed, convex.

Regarding QOT, it is immediate from Lemma 3 and 4 that QOT increases with an MPS of z also. The implication of this is that, if the DM is relatively more uncertain regarding social preferences for a particular wildland tract, then the marginal cost of its development is increased. In some sense, any particular message system becomes relatively more valuable when there is more uncertainty to be resolved and, by Theorem 3, more flexibility in current decisions is more valuable.

The Multi-Period Model

Extending the analysis of the previous two sections to a multi-period setting is straightforward. The results concerning QOT and the VOI were derived from a value function which could summarize optimal behavior in the second period of a two-period model or summarize optimal behavior over a longer time horizon.

The multi-period problem has the same structure as the two-period problem: in each time period the DM chooses the level of new development based on the amount of development up until that period and the DM's current probability function over z. The DM, after implementing q, observes the outcome of an experiment at the beginning of the next time period and updates his/her probability distribution over z using Bayes Rule. This new distribution forms the basis for the new choice of q.

Denote the probability distribution of z held by the DM at date t by F^t. Then

$$F^{t+1} = f(F^t, y^{t+1}),$$

where f is the map defined by Bayes Rule.

Let the conditional distribution of Y given the true value of the parameter z be $H(y|z)$. I assume in this section that z is stationary. Thus, the sequence of observations of y is a sequence of independently and identically distributed random variables drawn from the distribution H. Moreover, the sequence of distributions F^t forms a stationary Markov process. This structure allows me to call upon the basic mathematical results concerning stationary Markov decision problems. Toward this end, I introduce some terminology and notation.

Let the state of the system be $s\varepsilon S$, where $s=(D,F)$. Define $W(D,z) = U(D,1-D,z)$. Finally, let $g(D_t,F^t,q_{t+1},y^{t+1})$ be the transition equation, which gives tomorrow's state (D_{t+1},F^{t+1}) if today's state is (D_t,F^t), action q_{t+1} is implemented, and tomorrow's realization of the experiment is y^{t+1}.

The DM wishes to solve

$$\begin{aligned} &\max_{q_t} \quad E \sum_t a^t W(D_t,z) \\ &\text{s.t.} \quad D_{t+1} = D_t + q_{t+1};\ q_t \geq 0;\ D_t \varepsilon [0,1]. \end{aligned} \tag{32}$$

At each date the DM chooses q as a function of the current state, i.e., $q_{t+1}=q(D_t,F^t)$. The function q is called a plan. Given the plan q, define the expected discounted return by

$$I(q)\ (s) = E\ [W(D^o + q(D^o,F^o))] + \sum_{t=1}^{\infty} a^t\ W(D_{t-1} + q(D_{t-1},F^{t-1}),z)].$$

If Q is the set of all plans, the value function for the problem (32) is

$$V(s) = \sup_{q\varepsilon Q} [I(q)(s)].$$

Results concerning the decision problem (32) are stated with regard to the operator T defined by

$$TV(s) = \sup \{\int_Z W(D+q,z) + a \int_Y V(g(s,q,y)dH(y|z)\ dF(z)\}. \tag{33}$$

Naturally, $g(s,q,y) = (D+q,f(F,y))$.

If, in addition to the assumptions invoked so far in this chapter, it is assumed that U is bounded above, the transition probability

$$\text{Prob } (F^{t+1} | F^{t})$$

is weakly continuous, and Z and Y are finite, then it may be shown that T is a contraction map on a Banach space and, hence, has a unique fixed point, the value function for (32). Moreover, it may be concluded that the value function is a bounded, continuous, concave function, and that there exists a unique, stationary, continuous, optimal plan q(s) (Maitra, 1966; Easley and Spulber, 1982). Further, Blume et.al. (1982) show that the value function is differentiable.

Thus, if the functions V(•) and L*(•) are reinterpreted as value functions for the infinite horizon problem (32), all of the results derived above apply to solutions to that problem as well.

Discussion

So far, I have avoided any interpretation of the above model which would give it "policy relevance." In particular, I have avoided any statement of what constitutes experimentation. In this section, I offer some interpretive comments.

In many economic models of Bayesian learning either demand or technology (or both) is uncertain and experimentation consists of observations of current market outcomes. Some investigators of QOV have this type of process in mind in that experimentation consists of observing the level of benefits obtained from providing certain stocks of developed and preserved areas. This approach is quite clear in Arrow and Fisher (1974, p. 314).

However, if the multi-period, imperfect learning model considered here is to be non-trivial, this approach is unsatisfactory unless either: (1) the DM does not know (and cannot discover) the function form of U(•), or (2) the parameter z varies through time. For, if the form of U is known and the level of utility achieved can be observed, the parameter z can be calculated, and the model collapses to the two-period, perfect learning case.[9/] In this case, one naturally wonders what prevented this learning in any periods before t=0.

If z_t varies through time, even if the form of U(•) is known, information about past achieved levels of utility does not (necessarily) provide information about future preferences. However, if past realizations of z_t and future ones are correlated, then experimentation might consist of using past utility levels to uncover the process governing the movement of z_t. One plausible interpretation here is that there exists an unknown distribution from which the z_t are drawn and the DM's problem is to simultaneously estimate that distribution and choose the level of current investment. This type of model has been explored in the literature on Bayesian adaptive control. Note that, unlike the literature oriented toward analysis of market outcomes where the problem of simultaneity of price expectations formation and price formation must be confronted, I have avoided this here in positing the problem as one facing a central authority.

An important point raised by the economic literature in this area is that I have implicitly assumed that investment choices and the informativeness of the experiment are independent. Obviously, if the DM can produce information by choosing a larger investment, then the conclusion that less of an area is developed may not hold (see Prescott, 1972;

Grossman, Kihlstrom, and Mirman, 1977; and Freixas, 1981). An analysis of this issue is beyond the scope of this paper. However, it is worth pointing out that the two approaches to experimentation, the production of information approach and that considered here, are applicable in different decision-making situations.

Consider first the situation in which the development activity consists of energy exploration in potential wilderness.[10/] Obviously, this is an example of a case where development produces information about the value of more development. An analysis of this situation would require a re-examination of the production of information literature in light of the irreversibility aspect of wilderness development.[11/]

Consider next the situation in which the development activity has known benefits; for example, the harvest of (known) inventories of timber. In this case, most of the uncertainty in the problem surrounds the value of preservation, and it seems unlikely that harvesting more timber constitutes a more effective experiment concerning the benefits of providing some amount of wilderness.

A plausible interpretation of experimentation in this latter situation, that adopted in this paper, is that a central authority (e.g., a national Forest Service) is acting on behalf of its constituents to maximize community welfare. The central authority ignores distributional issues (or uses distributional weights) and conducts a benefit-cost type analysis of alternative preservation policies. However, since wilderness is a public good provided by the Forest Service, the benefits of providing it (the position of the compensated demands for it) are unknown. The experimenting authority tries to estimate the willingness-to-pay for wilderness by a variety of means. In particular, it might estimate a demand curve for recreational aspects of wilderness provision via a travel-cost approach, or undertake bidding-game analyses with a sample of its constituents, or hold public hearings. All of these constitute experimentation which provides some information to the agency about the benefits of providing wilderness.

Naturally, there are several questions raised by this interpretation that are not incorporated into the model above. Specifically, I have not addressed the menu of experiments that may be open to the DM and how (s)he might choose experiments over time (see, DeGroot, 1970). Further, I clearly have avoided any "bureaucratic behavior" interpretation of the function U(•) and how the preferences of individuals might relate to preferences of a government agency. And I have not discussed the strategic implications of experimenting via demand revealing processes (see, e.g., Dasgupta, 1982).

There are several further issues that might be addressed in future research. Suppose that wilderness decisions as well as development decisions are sufficiently costly to reverse as to be considered irreversible. For example, declassification of wilderness areas preserved by law might prove politically impossible. Then, a question arises as to the amount of information that should be obtained before a decision is implemented. In this sampling problem (and in the analysis of this paper as well) consideration should be made of the heterogeneity of the resource base and possibilities for substitution among wilderness areas both for wilderness services and the provision of commodities. One might seek, then, an idea of how wilderness decisions should be structured over time. Certainly, making all decisions at the initial date, or all decisions after an equal delay, will not be optimal.

FOOTNOTES

1/ Assistant Professor, Department of Agricultural and Applied Economics, University of Minnesota. Comments on an earlier draft by Lawrence Blume, Richard Porter, and participants in the Resource Economics Seminar at the University of Michigan are appreciated. In particular, I would like to thank Michael Moore for numerous conversations on these matters. Remaining errors and opacity are the responsibility of the author. Financial assistance from Resources for the Future is gratefully acknowledged.

2/ The concept was developed by Weisbrod (1964) and has been much discussed since. For recent summaries of the literature, see Bishop (1982), Graham (1981), and Smith (1983).

3/ Further confusion is provided by alternative definitions of the VOI (e.g., Lindley, 1971; Antonovitz and Roe, 1982), which resemble QOV, especially as the latter is presented by Henry (1974b). This issue is not explored here. It is worthwhile pointing out that the approach of Conrad (1981) is the same as that of Gould (1974).

4/ Compare Oi (1961), Sandmo (1971), and Hartman (1976). See Epstein (1978) for a complete analysis.

5/ Thus, I posit "rational expectations" in that the subjective distribution of the DM and the true distribution is the same.

6/ Note that in his proof of a similar proposition, Henry (1974b) implicitly imposes an assumption that $\bar{q}_1 = q_1^m$, without justification in the general case.

7/ One difference that arises is that in what follows I restrict attention to the finite or countable case, whereas in this section, I did not need to impose this restriction. Thus, there is not a perfect relationship between the two sections. However, a key result below (Lemma 2) has not been proven in this more general setting.

8/ Much of this material was completed (and presented to the Resource Economics Seminar at the University of Michigan) before the author became aware of the analysis of Epstein (1980). The exposition here follows this latter work; the substance of the two analyses was identical.

9/ More precisely, it is a multi-period model where the payoff for future periods is summarized by a $\{\max_q U(q,1-q,z^*\}/r$.

10/ For a discussion of this issue, see Runge (1983). Tiesberg (1980) considers exploration on public lands, but does not analyze the situation where the opportunity cost of exploration (foregone wilderness services) may be growing in value over time.

11/ The degree of irreversibility may not be absolute in this instance, and should even be considered a choice variable. This remark may apply to many possible wilderness development projects. Porter (1982) discusses this issue in a model without uncertainty.

REFERENCES

Antonovitz, F. and T. Roe, "A Measure of the Value of Information for the Competitive Firm Under Price Uncertainty," Staff Paper P82-28, Department of Agricultural and Applied Economics, University of Minnesota, 1982.

Arrow, K., "Optimal Capital Policy with Irreversible Investment," in J. Wolfe, ed., Value, Capital, and Growth, Edinburgh University Press, Edinburgh, 1968.

Arrow, K., and A. Fisher, "Environmental Preservation, Uncertainty, and Irreversibility," Quarterly Journal of Economics, 88:312-319, 1974.

Bernanke, B. S., "Irreversibility, Uncertainty, and Cyclical Investment," Quarterly Journal of Economics, 98:85-106, 1983.

Bishop, R., "Option Value: An Exposition and Extension," Land Economics, 58:1-15, 1982.

Blume, L., D. Easley, and M. O'Hara, "Characterization of Optimal Plans for Stochastic Dynamic Programs," Journal of Economic Theory, 28:221-234, 1982.

Conrad, R., "Quasi-Option Value and the Expected Value of Information," Quarterly Journal of Economics, 94:813-820, 1980.

Dasgupta, P., "Environmental Management Under Uncertainty," in Smith and Krutilla, eds., Explorations in Natural Resource Economics, Johns Hopkins Press, Baltimore, 1982.

DeGroot, M., Optimal Statistical Decisions, McGraw-Hill, New York, 1970.

Easley, D. and D. Spulber, "Stochastic Equilibrium and Optimality with Rolling Plans," International Economic Review, 22:79-103, 1981.

Epstein, L., "Production Flexibility and the Behavior of the Competitive Firm Under Price Uncertainty," Review of Economic Studies, 45:251-261, 1978.

Epstein, L., "Decision Making and the Temporal Resolution of Uncertainty," International Economic Review, 21:269-283, 1980.

Fisher, A., J. Krutilla, and C. Cicchetti, "The Economics of Environmental Preservation: A Theoretical and Empirical Analysis," American Economic Review, 62:605-619, 1972.

Freixas, X., "Optimal Growth with Experimentation," Journal of Economic Theory, 24:296-309, 1981.

Graham, D., "Cost-Benefit Analysis Under Uncertainty," American Economic Review, 71:715-725, 1981.

Grossman, S., R. Kihlstrom, and L. Mirman, "A Bayesian Approach to the Production of Information and Learning by Doing," Review of Economic Studies, 44:533-547, 1977.

Hanemann, W. M., "Information and the Concept of Option Value," Working Paper No. 228, Giannini Foundation of Agricultural Economics, University of California, 1983.

Hartman, R., "Factor Demand with Output Price Uncertainty," American Economic Review, 66:675-681, 1976.

Henry, C., "Investment Decisions Under Uncertainty: The 'Irreversibility Effect'," American Economic Review, 64:1006-1012, 1974a.

Henry, C., "Option Values in the Economics of Irreplaceable Assets," Review of Economic Studies: Symposium on the Economics of Exhaustible Resources, 89-104, 1974b.

Krutilla, J., "Conservation Reconsidered," American Economic Review, 47:777-786, 1967.

Krutilla, J., Natural Environments: Studies in Theoretical and Applied Analysis, Johns Hopkins Press, Baltimore, 1972.

Krutilla, J. and A. Fisher, The Economics of Natural Environments, Johns Hopkins Press, Baltimore, 1975.

Lindley, D., Making Decisions, Wiley-Interscience, London, 1971.

Maitra, A., "Discounted Dynamic Programming on Compact Metric Spaces," Sankhya, Series A., 30, Part 2:211-216, 1968.

Marschak, J. and K. Miyasawa, "Economic Comparability of Information Systems," International Economic Review, 9:137-174, 1968.

Oi, W., "The Desirability of Price Instability Under Perfect Competition," Econometrica, 29:58-64, 1961.

Porter, R., "The New Approach to Wilderness Preservation through Benefit-Cost Analysis," Journal of Environmental Economics and Management, 9:59-80, 1982.

Porter, R. C., "The Optimal Timing of an Exhaustible, Reversible Wilderness Development Project," Institute of Public Policy Studies, Discussion Paper 178, University of Michigan, 1982.

Prescott, E., "The Multi-Period Control Problem Under Uncertainty," Econometrica, 6:1043-1058, 1972.

Rothschild, M. and J. Stiglitz, "Increasing Risk I: A Definition," Journal of Economic Theory, 2:225-243, 1970.

Runge, C. F., "Energy Exploration on Wilderness: A Note on 'Privitization' and Public Lands Management," Unpublished Manuscript, Department of Agricultural and Applied Economics, University of Minnesota, 1983.

Sandmo, A., "On the Theory of the Competitive Firm Under Price Uncertainty," American Economic Review, 61:65-73, 1971.

Smith, V. K., "Option Value: A Conceptual Overview," Southern Economic Journal, 49:654-668, 1983.

Tiesberg, T., "Federal Management of Energy and Mineral Resources on Public Lands," Bell Journal of Economics, 11:448-465, 1980.

Weisbrod, B., "Collective-Consumption Services of Individual-Consumption Goods," Quarterly Journal of Economics, 78:471-477, 1964.