DISCUSSION PAPER NO. 197

STOCHASTIC EVOLUTION AND CONTROL OF AN ECONOMIC ACTIVITY

by

S. D. Deshmukh* and S.D. Chikte+

February 1976

This research was supported by NSF Grant GS-28287

*Department of Managerial Economics and Decision Sciences,
Graduate School of Management
Northwestern University
Evanston, Illinois 60201

+Systems Analysis group
Department of Electrical Engineering
Polytechnic Institute of New York
Brooklyn, New York

ABSTRACT

An economic activity, whose progress status can be assessed during its conduct, evolves stochastically as a result of various uncertainties that are only partially controlled by resource expenditures. Suppose that the reward from pursuing the activity is collected only upon its termination and depends on its final status. Then the problem of optimally controlling the activity involves determination of a dynamic resource allocation strategy as well as a stopping rule for terminating the activity. Thus, the problems of selecting an optimum goal that the activity must pursue and that of determining the optimum strategy for attaining that goal are considered in an integrated fashion. The optimal solutions to these problems are shown to have economically meaningful characterizations under reasonable assumptions.

STOCHASTIC EVOLUTION AND CONTROL OF AN ECONOMIC ACTIVITY

1. INTRODUCTION

Consider the following general framework encompassing a wide class of problems of dynamic resource allocation to an economic activity that is to be pursued under uncertainty. We are given an activity to which effort is allocated over a period of time, and in which it is meaningful and possible to assess its status at any time during the course of its evolution. The effect of an allocation of effort on the status of the activity is stochastic due to internal and external uncertainties. Allocating a large amount of effort costs more but also yields greater progress on average, (resulting in an improvement in the activity status), though perhaps at a decreasing marginal rate. On the other hand, allocating too small an effort, though economical in the short run, may cause a setback (resulting in a deterioration of the activity status) due to stochastic obsolescence and undesirable environmental disturbances. Suppose that once the allocation of effort is stopped and the activity is terminated, a reward is collected, whose value depends upon the terminal status of that activity. At any point in time during the course of such an activity, the manager in charge faces two interrelated problems: (1) whether he should allocate any further expenditure of resources to the activity or he should conclude its pursuit and collect the terminal reward, and (2) in case he decides to continue the activity, determine the amount of further resource allocation. Since the resources are expended during the course of the activity with the expectation of obtaining a terminal return, such an activity will be termed a "project" throughout this paper.

The above framework may be viewed as an attempt at combining the salient features of the optimal stopping problem (e.g. see Breiman [3], and Chow, Robbins and Siegmund [3]) and the optimal first passage (or optimal pursuit) problem (e.g. see Derman [6] and Eaton and Zadeh [8]). Thus, in addition to the usual stop or continue decisions in the optimal stopping problem, we allow for the intensity of continuation to be a decision variable as well. Similarly, rather than optimally pursuing a given fixed target, as in the first passage problem, we allow for the determination of the target itself to be a decision variable along with the optimal strategy for pursuing it.

As an example of evolution and control of such an activity consider a product development project, whose status at any intermediate point during its conduct may be represented by the quality ranking of the product developed so far, relative to other similar products available in the market at that time. As in Dorfman and Steiner [6], this relative quality may be looked upon as a shift parameter in the demand function of that product. A better product at the same production cost and selling price yields greater sales and higher profits. This relative product quality may thus be measured in terms of the potential profit obtainable from marketing the product. During the product development, its relative quality changes stochastically due to the internal technological uncertainties (inherent in the R and D process) and the external market uncertainties (due to random changes in consumer tastes and competing products affecting the product demand). The development expenditures partially control the technological uncertainty. Once the development project is terminated, the final product is introduced into the market, thereby yielding the net expected discounted profit from then on as the terminal reward. During the conduct of the development project the R and D manager must decide (1) whether it is worthwhile improving further the quality of the product developed so far, rather than marketing it as it is, and (2) if so, how to finance the further product development through time, taking into account the above mentioned uncertainties. This model allows for a degree of partial success of the project along with the freedom to select the target product quality to be developed. In the standard literature on R and D project funding over time (e.g. see Aldrich and Morton [1], Gittins [9], Hess [10], Kamien and Schwartz [11], and Lucas [13]), on the other hand, the state of the project at any time is usually classified as being either completely successful or not, where "success" is prespecified (corresponding to a fixed target or goal) and a partial success during the progress of the project is not only meaningless but also inconsequential as far as the funding strategy is concerned.

As another example, consider the activity of searching for a job, (or a house or a secretary). The intermediate status of such a project may be summarized by the best offer on hand at that time. As the search process continues, the project status may improve due to an arrival of a better offer or it may deteriorate due to losing the currently best offer to a competitor. Both events are uncertain, although the probability of an improvement can be increased by increasing the intensity of search. Thus the currently best offer is not always lost (as in sampling without recall, DeGroot [5], ch.13) nor is it always available in the next period (as in sampling with recall), but improves or decays stochastically and may be partially controlled by expenditure of search effort. Given the best offer on hand so far, the searcher must decide (1) whether to accept this offer (which then corresponds to the terminal reward) and end the search or to continue the search for a better offer at the risk of losing the one on hand, and (2) in case of continuation, determine the intensity with which he should search. Such a formulation generalizes the usual search problem (e.g. see DeGroot [5], ch.13, Kohn and Shavell [12], and McCall [15]) so that, in addition to determining the optimal minimum acceptable offer, a strategy for optimally allocating the search effort to achieve it must also be selected.

The model presented in the next section attempts to capture the essence of the above framework and is similar to the controlled random walk model proposed by Radner and Rothschild [16] for describing the progress status of an activity. (They, however, consider the implications of following various plausible, but not necessarily optimal, behavioral rules for distributing resources among several activities.) In the third section, we establish the structure of the optimal stopping region and that of the optimal resource allocation strategy in the continuation region by applying the methodology of Markov decision processes. The final section concludes with some remarks and their further implications.

2. A STOCHASTIC EVOLUTION MODEL

It is often possible to evaluate the status of a project at any time during the course of its conduct in terms of the performance level attained so far. As in the illustrative examples of the previous section we will assume that the performance level of the project is reviewed at discrete points in time and at any date n it is measured by the potential reward that could be collected from commissioning the project in the current state, to be denoted by x_n, $n=0,1,2,\ldots$. Without loss of generality, we may take $x_n \in [0,1]$, where 1 denoted the ideal performance level and the maximum terminal reward attainable with the limited technological capability of the capital and labor resources on hand, while 0 corresponds to the worst possible project performance that is completely worthless.

The project evolution over time consists of a series of intermediate successes and setbacks, as summarized by the sequence of performance levels $\{x_n:n=0,1,2,\ldots\}$ perceived by the manager. The project performance level changes through time as a result of the internal technological uncertainties, the external environmental uncertainties and the expenditure of resources counteracting the undesirable effects of both the kinds of uncertainties.

At any review point n the manager may decide to terminate the project (because, perhaps, he deems the currently attained performance level x_n satisfactory enough) and collect the reward x_n or he may decide to continue the project (in the hope of perhaps improving the project status before commissioning it). If he decides to continue, he must select a resource allocation, which may be aggregated into the monetary expenditure $b_n \in [0,B]$, where $B > 0$ corresponds to the upper limit on the available resources at the manager's disposal. Thus, at any stage n, the set of available actions may be denoted by $A = \{s,c\} \times [0,B]$, where s and c correspond to stopping and continuing actions, respectively. Clearly, upon taking the termination action s no further allocation is necessary. The terminal reward as well as the interim expenditures are discounted at rate $\beta \in [0,1)$, i.e. β is the present value of one dollar in the next period.

An allocation b may result either in an improvement in the performance level or in its deterioration, the actual outcome being uncertain due to technological and environmental factors. Let $p(b)$ denote the probability of improvement, (i.e. $[1-p(b)]$ is the probability of deterioration). We will assume $p(\cdot)$ to be a continuous nondecreasing function on [0,B] with $p(0) = 0$. Thus, a positive allocation is necessary to achieve progress and a greater allocation is more likely to be successful in counteracting the unfavorable elements in the environment than a smaller one.

At any stage, if the allocation is successful, the magnitude of improvement in the performance level may depend upon the current performance level x, and will be denoted by $U(x)$, where $U(\cdot)$ is taken to be a continuous function on [0,1]. Similarly, the magnitude of possible deterioration due to an unsuccessful allocation will be denoted by $L(x)$, if the current performance level is x, $L(\cdot)$ being continuous on [0,1]. Thus, the project performance level fluctuations may be modelled by the following controlled Markov process. For any $x_n \in [0,1]$, $b_n \in [0,B]$, $a_n = (c,b_n)$, $n=0,1,2,\ldots$

$$(2.1) \qquad x_{n+1} = \begin{cases} x_n + U(x_n) & \text{with probability } p(b_n) \\ x_n - L(x_n) & \text{with probability } [1-p(b_n)]. \end{cases}$$

For technical reasons, we will assume that once a stop action (i.e. $a_n = (s,b_n)$) is taken, the process goes to an arbitrary state, say $-\infty$, and remains there then on, incurring no additional costs and yielding no further reward. If an allocation is successful, the resulting performance level will be assumed to be higher if the current performance level itself is higher, i.e. $x + U(x)$ is nondecreasing in x. Analogously, in case of a deterioration due to an unsuccessful allocation, the lower the current performance level the lower is the reduced performance level, i.e. $x - L(x)$ is nondecreasing in x. Thus, starting a stage with a higher performance level is always better. Finally, we may take $U(1) = L(0) = 0$, implying the impossibility of an improvement upon the ideal and a deterioration below the worst level. These conditions on the functions $p(\cdot)$, $U(\cdot)$ and $L(\cdot)$ may be summarized as

Assumption 1: $p(0) = 0$, $p(b_2) \geq p(b_1)$, if $B \geq b_2 \geq b_1 \geq 0$;
$U(x) \geq 0$, $L(x) \geq 0$ if $x \in [0,1]$; $U(1) = L(0) = 0$;
$x_2 + U(x_2) \geq x_1 + U(x_1)$, $x_2 - L(x_2) \geq x_1 - L(x_1)$, if $1 \geq x_2 \geq x_1 \geq 0$.

Note that Assumption 1 ensures that the dynamics (2.1) is well defined, i.e. $x_{n+1} \in [0,1]$ whenever $x_n \in [0,1]$.

The project manager controls the evolution of such a project by employing a strategy Π, which, in general, is a sequence of (possibly randomized) decision rules $\{\Pi_n : n=0,1,2,...\}$, that specifies an action $a_n \in \{s,c\} \times [0,B]$ at each date n, as a (Borel measurable) function of its previous history $(x_0, a_0, ..., x_n)$. A stationary strategy is a particularly simple type of strategy given by a (Borel measurable) function $\alpha : [0,1] \to \{s,c\} \times [0,B]$, so that on any date, if the current performance level is x, the action $\alpha(x)$ is specified by the strategy. Such a strategy is attractive by virtue of its modest informational requirements as well as the ease of its implementation.

Starting with an initial performance level x_0 and following a strategy Π, let

(2.2) $N = N(x_0,\Pi) = \text{Inf} \{n : a_n \in \{s\} \times [0,B]\}$

be the random stopping time (possibly infinite valued) at which the project is terminated and the reward x_N collected. The net expected discounted return, starting in x_0 and following Π may then be denoted by

(2.3) $$W(x_0,\Pi) = E[\beta^N x_N - \sum_{n=0}^{N-1} \beta^n b_n | x_0,\Pi].$$

Finally, let

(2.4) $$V(x_0) = \sup_{\Pi} W(x_0,\Pi), \quad x_0 \in [0,1]$$

to be called the Project Value Function, i.e. $V(x_0)$ is the net expected discounted return from the project, starting in state x_0 and optimally pursuing its further development. A strategy Π^* is said to be optimal at x_0 if $W(x_0,\Pi^*) = V(x_0)$ and it is said optimal if it is optimal at each $x_0 \in [0,1]$.

The objective of our project manager is to select an optimal strategy. Such a strategy will determine whether it is worthwhile improving the current project status, and, if so, how best to finance it, depending on its progress, and, finally, at what point N it is best to terminate it and collect x_N. This problem can be naturally analyzed using the theory of Markovian decision processes with discounting (see, for example, Blackwell [2], Strauch [18], and Ross [17],ch.6).

We establish the existence of a (measurable) optimal stationary strategy α^* and the optimality equation satisfied by the (measurable) project value function $V(\cdot)$ by invoking the results of Maitra [14] upon verifying that his topological and measure theoretic conditions hold in our problem. Thus the action space $A = \{s,c\} \times [0,B]$ is a compact metric space, while the immediate return function $f : [0,1] \times \{s,c\} \times [0,B] \to \mathbb{R}$ given by

$$(2.5) \qquad f(x,a) = \begin{cases} -b & \text{if } a = (c,b), \\ x-b & \text{if } a = (s,b) \end{cases}$$

is bounded and upper semi-continuous, as required. Also, if $h(\circ)$ is any bounded continuous function on $[0,1]$, then let $E[h|x,a] = h(x+U(x))\, p(b) + h(x-L(x))[1-p(b)]$ if $a = (c,b)$ and $E[h|x,a] = h(-\infty)$ if $a = (s,b)$, so that, by continuity of $p(\cdot)$, $U(\cdot)$ and $L(\cdot)$ it follows that $E[h|x_i,a_i] \to E[h|x,a]$, whenever $x_i \to x$ and $a_i \to a$, thus verifying the weak continuity of the transition law. Hence, an application of Maitra's main theorem (based on a selection theorem of Dubins and Savage) to our problem immediately yields the following

Proposition 1. The project value function $V : [0,1] \to \mathbb{R}$ is upper-semicontinuous (hence Borel measurable) and satisfies the following optimality equation

(2.6) $V(x) = \text{Max}\{x, Q(x)\}, \; x \in [0,1]$, where

(2.7) $$Q(x) = \sup_{b \in [0,B]} \{-b + \beta\, p(b)\, V(x+U(x)) + \beta[1-p(b)]V(x-L(x))\}$$

and there exists a (measurable) optimal stationary strategy $\alpha^*: [0,1] \to \{s,c\} \times [0,B]$, which, when in state x, specifies stopping if $x \geq Q(x)$ and continuation otherwise; along with an allocation attaining the supremum in (2.7).

The function $Q(\cdot)$ of (2.7) gives the optimal return if we are forced to continue for one stage and if we follow an optimal strategy then on. It will be found convenient to decompose the strategy α^* into the <u>optimal stopping strategy</u> $\alpha_1^* : [0,1] \to \{s,c\}$ for determining the project termination and the <u>optimal allocation strategy</u> $\alpha_2^* : [0,1] \to [0,B]$ for specifying a resource allocation in case of continuation, so that $\alpha^* = (\alpha_1^*, \alpha_2^*)$. In case there are more than one actions attaining the maximum in (2.6) and (2.7), α^* will be assumed to specify stopping rather than continuation and less expenditure rather than more, thus implying a risk averting and thrifty project manager.

In the next section, we investigate the structural properties of optimal strategies $\alpha_1^*(\cdot)$ and $\alpha_2^*(\cdot)$ and the project value function $V(\cdot)$.

3. OPTIMAL CONTROL STRATEGY

In order to characterize the form of the strategy $\alpha^*(\cdot)$ and the value function $V(\cdot)$ we will need to make further assumptions. The following assumption is regarding the magnitudes of possible changes in the project status.

Assumption 2: $U(\cdot)$ is convex and $L(\cdot)$ is concave on $[0,1]$, i.e. if x_1, x_2, $\lambda \in [0,1]$, then

$$U(\lambda x_1 + (1-\lambda)x_2) \leq \lambda\, U(x_1) + (1-\lambda)U(x_2) \text{ and}$$

$$L(\lambda x_1 + (1-\lambda)x_2) \geq \lambda\, L(x_1), + (1-\lambda)L(x_2).$$

This assumption, together with $U(1) = L(0) = 0$, implies that $U(\cdot)$ and $L(\cdot)$ are continuous on $[0,1]$ and monotonically decrease to 0 as $x \to 1$ and $x \to 0$, respectively. Thus, an already high performance level can not be further improved upon significantly due to saturation effects and, furthermore, these effects become more accentuated as the performance level approaches the ideal level 1 attainable with the available technology. Equivalently, as the return obtainable from the project of current status increases, the additional return obtainable from a further improvement in the project status diminishes, though at a diminishing rate. This corresponds to the usual assumption of diminishing marginal returns in economic analysis. Similarly, given an already poor performance level, there is little more to lose in spite of an unsuccessful allocation and, as the current performance level improves, although the deterioration, if it occurs, increases, it does so at a diminishing rate.

The following proposition summarizes certain properties of the project value function that will be needed later.

Proposition 2: Under the Assumptions 1 and 2 the project value function $V(\cdot)$ satisfies

(3.1) $1 \geq V(x) \geq 0$ for all $x \in [0,1]$ (nonnegativity)

(3.2) $V(x_1) \geq V(x_2)$ if $1 \geq x_1 \geq x_2 \geq 0$ (monotonicity)

(3.3) $\lambda V(x_1) + (1-\lambda)V(x_2) \geq V(\lambda x_1 + (1-\lambda)x_2)$, for all λ, x_1, $x_2 \in [0,1]$ (convexity).

Proof: Here (3.1) follows directly from (2.6). To show (3.2) and (3.3), consider the sequence of functions $\{V_n\}_{n=0}^{\infty}$ on $[0,1]$ defined recursively by,

$$V_0(x) = x$$

(3.4) $$V_{n+1}(x) = \text{Max}\{x, Q_{n+1}(x)\}, \text{ where}$$

$$Q_{n+1}(x) = \sup_{b \in [0,B]} \{-b + \beta\, p(b)\, V_n(x+U(x)) + \beta[1-p(b)]\, V_n(x-L(x))\}$$

Then it can be shown (as e.g. in Ross [17] ch.6) that $\lim_{n\to\infty} V_n(x) = V(x)$

uniformly in $x \in [0,1]$, so that it sufficies to show monotonocity and convexity of each $V_n(\cdot)$. Clearly $V_0(x)$ has these properties. Suppose that $V_n(\cdot)$ also has these properties. Since $x + U(x)$ and $x - L(x)$ are nondecreasing and convex in x by assumptions 1 and 2, it follows that $V_n(x+U(x))$ and $V_n(x-L(x))$ are nondecreasing and convex in x, so that $Q_{n+1}(x)$ is nondecreasing and convex. From (3.4) it follows that $V_{n+1}(x)$ is nondecreasing and convex, completing the induction argument.

Q.E.D.

An immediate consequence of (2.7), (3.1), (3.2), and (3.3) is

Corollary: The function $Q(\cdot)$ is nonnegative, bounded by 1, nondecreasing and convex.

Thus, optimal project management never results in losses, while taking up the project in a better initial state and continuing optimally will yield a higher net return on average, and, moreover, this relative advantage due to starting with a better project state increases as we get closer to the ideal level.

These properties of the project value function can now be used to characterize the optimal stopping strategy $\alpha_1^*(\cdot)$.

Proposition 3: Under the Assumptions 1 and 2 there exists $x^* \in [0,1]$ such that

$$\alpha_1^*(x) = s \qquad \text{if and only if} \quad x \in [x^*,1].$$

Proof: By nonnegativity of $V(\cdot)$, either $V(0) = 0$ or $V(0) > 0$ are the only two possible cases.

Case 1: If $V(0) = 0$, then $Q(0) = 0$. Monotonicity and convexity of Q, together with $Q(x) \leq 1$ for all $x \in [0,1]$, implies that $Q(x) \leq x$ for all $x \in [0,1]$. Hence, $\alpha_1^* \equiv s$ and we define $x^* = 0$.

Case 2. If $V(0) > 0$, then $Q(0) > 0$, moreover $Q(0) < 1$ since $p(0) = 0$. Now consider the set

$$T = \{x \in [0,1]: \ Q(x) = x\}.$$

We claim that either T is an interval $[x^*,1]$ or that $T = \{x^*\}$ or $T = \{x^*,1\}$ for xome $x^* < 1$ and in each case, $Q(x) \leq x$ for all $x < x^*$, thereby proving the proposition.

We first show that $T \neq \{1\}$. Otherwise, since $U(x)$ is convex and monotone decreasing to 0 as $x \to 1$, it is clear that $N(x,\alpha^*) = \infty$ with probability 1, for all $x < 1$, so that from (2.3) and (2.4) $V(x) = W(x,\alpha^*) \leq 0$, implying that $V(x) = 0$ for all $x \in [0,1)$, a contradiction, since,in case 2,$V(x) > 0$ for all $x \in [0,1]$.

Next, suppose that T is countable, i.e. $T = \{x_n^*,\ n \geq 2$: all x_n^*'s are distinct$\}$. Now, by convexity of $Q(\cdot)$ we must have $n \leq 2$. If $T = \{x_1^*,x_2^*\}$ with $1 > x_2^* > x_1^* > 0$, then again by convexity of Q, we must have $Q(x) \geq x$ if $x \in [0,x_1^*]$, $Q(x) < x$ if $x \in (x_1^*,x_2^*)$ and $Q(x) \geq x$ if $x \in [x_2^*,1]$. But then, if we stop in $x \in [x_2^*,1]$ the terminal reward is at least x_2^*, while α_1^* specified continuation, so that following α_1^* we stop only in a state $x \in (x_1^*,x_2^*)$, receiving less than x_2^*. This contradicts optimality of α_1^*. Hence, if T contains discrete points then it must be a singleton set $\{x^*\}$ or $\{x^*,1\}$ and moreover $x^* < 1$.

By monotonicity and convexity of Q and that $Q(1) \leq 1$, the only other possibility is that $T = [x^*,1]$ for some $x^* < 1$, so that $Q(x) \geq x$ for all $x \in [0,x^*]$ and $Q(x) \leq x$ for all $x \in [x^*,1]$.

Q.E.D.

Thus, there exists a critical performance level x^* that the manager should attempt to attain before commissioning the project and collecting the reward. If the project performance level exceeds the critical level, then a further expenditure of resources is not worthwhile in relation to the additional potential benefits that may accrue. It is interesting to note that if the optimal return from the worst performance level V(0) is zero, then regardless of the actual initial performance level it is best not to attempt to improve it at additional expenditures.

In order to characterize the structure of the optimal resource allocation strategy $\alpha_2^*(\cdot)$ in the continuation region $[0,x^*)$ we need the following

Assumption 3: $U(x) + L(x)$ is monotone, nondecreasing in $x \in [0,1]$ i.e. if $1 \geq x_2 \geq x_1 \geq 0$, then

$$U(x_2) + L(x_2) \geq U(x_1) + L(x_1).$$

In order to interpret this assumption we may write $U(x) + L(x)$ as $[x+U(x)] - [x-L(x)]$, which is the opportunity loss in the performance level due to suffering a setback rather than a success in one stage, starting with level x. Then Assumption 3 says that, this one stage opportunity loss increases as the current project performance level approaches the ideal one, so that the nearer we are to the ideal level, the more the success in each stage of the project counts as against a setback.

We are now ready to state

Proposition 4: Under the Assumptions 1, 2 and 3, the optimal resource allocation strategy α_2^* is a strictly positive nondecreasing function on $[0,x^*)$ and is identically 0 on $[x^*,1]$, where x^* is as in Proposition 3.

Proof: Assume that $x^* > 0$, for otherwise the Proposition follows vacuously. Clearly, $\alpha_1^*(x) = s$ implies $\alpha_2^*(x) = 0$ and the entire process stops. If $x \in [0,x^*)$, then $\alpha_1^*(x) = c$ and $V(x) = Q(x)$, so that, from Proposition 1, $\alpha_2^*(x)$ attains the supremum in

$$(3.5) \quad V(x) = \sup_{b\ [0,B]} \{F(x,b)\}, \text{ where}$$

$$(3.6) \quad F(x,b) = -b + \beta p(b)\ V(x+U(x)) + \beta[1-p(b)]\ V(x-L(x)),\ x \in [0,x^*)$$

Suppose $\alpha_2^*(x) = 0$ for some $x \in [0,x^*)$. Then, from (3.5), (3.6) and the fact that $p(0) = 0$ we have $V(x) = \beta\ V(x-L(x))$, implying $V(x) = 0$, since $V(x) \geq 0$, $L(x) \geq 0$ and $V(\cdot)$ is nondecreasing. But $V(x) > x \geq 0$, for all $x \in [0,x^*)$, yielding a contradiction. Hence $\alpha_2^*(x) > 0$ for all $x \in [0,x^*)$.

To show that $\alpha_2^*(x_2) \geq \alpha_2^*(x_1)$ whenever $x^* > x_2 \geq x_1 \geq 0$ it suffices to show that $[F(x_2,b) - F(x_1,b)]$ is nondecreasing in $b \in [0,B]$. For otherwise we get

$$F(x_2,\alpha_2^*(x_1))-F(x_1,\alpha_2^*(x_1)) > F(x_2,\alpha_2^*(x_2))-F(x_1,\alpha_2^*(x_2))$$

contradicting optimality of α_2^*. That $[F(x_2,b)-F(x_1,b)]$ is non-decreasing in b follows from the assumptions and convexity and monotonicity of $V(\cdot)$. Q.E.D.

Thus, in the continuation region, the optimal allocation strategy is never passive and the closer we get to the optimum goal x^* the more optimistic our outlook becomes and the more we are encouraged to strive harder to attain the goal, implying the desirability of an aggressive and vigorous management.

4. CONCLUSION

In section 2 we proved the existence of a stationary optimal project management strategy α^*, while in section 3 we established its economically meaningful characterization, which may now be summarized as in (4.1) below.

$$(4.1) \qquad \alpha^*(x) = \begin{cases} (s,0) & \text{if } x \in [x^*,1] \\ (c,\alpha_2^*(x)) & \text{if } x \in [0,x^*), \end{cases}$$

where the target critical level is strictly less than 1 and is given by

(4.2) $x^{*} = \mathrm{Inf}\{x \in [0,1] : Q(x) = x\}$

and the allocation strategy $\alpha_2^*(\cdot)$ is strictly positive and nondecreasing over the continuation region $[0,x^*)$. Thus, as in the optimal stopping problem, the manager selects the optimum goal x^*, so that the project terminates as soon as the goal is attained or exceeded. Then, given such a goal x^*, the manager determines a resource allocation strategy $\alpha_2^*(\cdot)$ for optimally attaining it, as in the first passage (or optimal pursuit) problem. The above characterization also enables us to decompose (2.4) as

(4.3) $V(x_0) = \sup_{x\in[0,1]} \sup_{\alpha_2\in\Delta_x} W(x_0,(x,\alpha_2)),$

where (x,α_2) denotes the policy that specifies stopping upon reaching x and Δ_x is the class of all stationary nondecreasing allocation strategies on $[0,x)$. Thus the entire project control procedure may be roughly described as follows. Select a goal x that the manager should strive for and determine an optimal allocation strategy $\alpha_2^x(\cdot)$ for attaining the goal (using, say, a policy improvement routine), yielding $W(x_0,(x,\alpha_2^x))$. Now "appropriately" modify the goal and repeat the process until a "satisfactory" combination of the goal and the corresponding resource allocation strategy for attaining it is obtained. Such a heuristic procedure has an obvious behavioral interpretation and may serve as a basis for developing a convergent algorithm under more specific assumptions such as finite state and action spaces.

REFERENCES

[1] Aldrich, C. and T. E. Morton, "Optimal Funding Paths for a Class of Risky R & D Projects," Management Science, Vol. 21, No. 5, pp. 491-500 (1975).

[2] Blackwell, D., "Discounted Dynamic Programming," Annals of Mathematical Statistics, Vol. 36, pp. 226-235, (1965).

[3] Breiman, L., "Stopping Rule Problems," Ch. 10, in Applied Combinatorial Mathematics, E. F. Beckenbach (Editor),

[4] Chow, Y. S., H. Robbins and D. Siegmund, Great Expectations: The Theory of Optimal Stopping, Houghton-Mifflin, 1971.

[5] DeGroot, M. H., Optimal Statistical Decisions, McGraw-Hill, (1970).

[6] Derman, C., Finite State Markovian Decision Processes, Academic Press, 1970.

[7] Dorfman R., and P. O. Steiner, "Optimal Advertising and Optimal Quality," American Economic Review, pp. 826-836, (1954).

[8] Eaton, J. H., and L. A. Zadeh, "Optimal Pursuit Strategies in Discrete State Probabilistic Systems," Trans. ASME Ser. D., J. Basic Engineering, Vol. 84, pp. 23-29 (1962).

[9] Gittins, J. C., "Optimal Resource Allocation in Chemical Research, "Advances in Applied Probability, Vol. 1, pp. 238-270 (1969).

[10] Hess, S. W., "A Dynamic Programming Approach to R&D Budgeting and Project Selection," IRE Transactions on Engineering Management, Vol. EM 9, pp. 170-178 (1962).

[11] Kamien, M. I. and N. L. Schwartz, "Expenditure Patterns for Risky R&D Projects," Journal of Applied Probability, Vol. 8, pp. 60-72 (1971).

[12] Kohn, M. G., and S. Shavell, "The Theory of Search," Journal of Economic Theory, June 1974.

[13] Lucas, R. E., "Optimal Management of an R&D Project," Management Science, Vol. 17, pp. A679-A697, (1971).

[14] Maitra, A., "Discounted Dynamic Programming on Compact Metric Spaces," Sankhya, series A. Vol. 30, part 2, pp. 211-216, (1968).

[15] McCall, J. J., "Economics of Information and Job Search," Quarterly Journal of Economics, Feb. 1970.

[16] Radner, R., and M. Rothschild, "On the Allocation of Effort," Journal of Economic Theory, June 1975.

[17] Ross, S. M., Applied Probability Models with Optimization Applications, Holden Day, (1970).

[18] Strauch, R. E., "Negative Dynamic Programming," Annals of Mathematical Statistics, Vol. 37, pp. 871-890 (1966).