THE ECONOMICS OF INFORMATION AND OPTIMAL STOPPING RULES

John J. McCall

October 1964

P-2987

BLANK PAGE

THE ECONOMICS OF INFORMATION AND OPTIMAL STOPPING RULES*

John J. McCall

The RAND Corporation, Santa Monica, California

1. INTRODUCTION

Recent work has emphasized the importance of information in a variety of economic problems.(9,14) Previously, the role of information in economics, while recognized as significant, was never analyzed. Methods now exist that permit a fairly precise evaluation of information for many important decision problems.

In this paper, two different stopping environments are analyzed. The decision process occurring in the first environment is continuous and repetitive like inventory management, process control, and maintenance management. However, the probability mechanism associated with the repetitive process is not known. Frequently, the decision-maker has the option of accumulating information about the probability law. The problem is when to stop collecting information. Statistical decision theory will be used to solve this class of problems.

In the second type of stopping environment the essence of the decision process is that it eventually terminates. Examples of such processes include house-searching, mate selection, and many kinds of investment decisions. It is assumed that the probability law governing

---

*Any views expressed in this paper are those of the author. They should not be interpreted as reflecting the views of The RAND Corporation or the official opinion or policy of any of its governmental or private research sponsors. Papers are reproduced by The RAND Corporation as a courtesy to members of its staff.

The paper was written while the author was a National Science Foundation postdoctoral fellow at the International Center for Management Science of the Netherlands School of Economics.

these terminating processes is known. Dynamic programming will be used to solve this class of problems.

The two environments are similar in that the stopping decision is in both cases dominated by economic considerations.

The maintenance of stochastically failing equipment constitutes a repetitive decision process and exemplifies the first class of stopping rules. In almost all practical applications of maintenance scheduling the stochastic law governing the equipment's failure behavior is not known perfectly. However, knowledge of the stochastic law can be improved by implementing a test program of one kind or another. The appropriate amount of information to purchase depends, of course, on the value of information relative to its acquisition cost. Usually, acquisition cost is easily estimated, and estimating the value of information is a simple application of statistical decision theory.

Investment or capital budgeting problems illustrate the second class of stopping problems. The investment problems considered here have the following structure. The decision-maker moves from one investment opportunity to another according to a discrete parameter Markov chain with known transition probabilities. In addition, the rate of return associated with each investment opportunity is a random variable with a known probability distribution--each state of the Markov chain is characterized by a probability distribution of rates of return. Finally, the time between investment opportunities is also a random variable with a known probability distribution. To recapitulate, the movement from investment opportunity i to investment opportunity j is determined by the transition matrix (P_{ij}); however, the movement from i to j is not simultaneous but is delayed, the length of the delay being given by a draw from the delay distribution; and entry into state j is accompanied by a draw from the rate of return distribution associated with state j. This stochastic process is a generalized semi-Markov process. Semi-Markov processes were first introduced by Levy and Smith* and have been used to analyze a variety of problems in applied probability. A semi-Markov process differs from the process described here in that the opportunity or rate of return distribution is degenerate and concentrated on a single rate of return. The extension of semi-Markov processes to include non-degenerate state distributions was first introduced by Karlin.[7]

Two different kinds of costs are present in this process. If the decision-maker decides to wait for another investment opportunity, then the investment is postponed in accordance with the delay distribution and a discounting fee is incurred. The second cost is a fixed fee charged whenever an investment is drawn from the rate of return distribution. For specified costs and probability distributions, the decision problem is whether to accept a known investment opportunity or collect more information and wait until at least one more opportunity is uncovered. As before, the investment decision depends on the cost of additional information relative to its expected value.

The stochastic investment process is analyzed for two different modes of accumulating investment opportunities. First, the process is assumed to have no memory. The rejection of any investment opportunity

---

*For a description of the semi-Markov process together with its history, see Ref. 12.

is permanent and the decision is between accepting the currently available opportunity or waiting for at least one more opportunity; secondly, the process is assumed to remember the highest rate of return observed. The decision-maker must then choose between investing in the maximum known rate of return or waiting for at least one more opportunity in hopes that it will exceed this maximum. The theory of optimal stopping[1,6,8] together with recent work on renewal programming[3] provide the methods that are used to calculate the value of information in these stochastic investment problems.

2. THE COST AND VALUE OF INFORMATION IN A SIMPLE SCHEDULING PROBLEM

2.1 Introduction

The purpose of this section is to apply some recent results of statistical decision theory to a particular replacement problem. It is hoped that the application will illustrate the relevance of decision theory to the problem of determining the appropriate scope and level of data collection activities.

As will be shown the value of information to a particular decision process can be measured prior to its purchase; the measure is in terms of the improved performance of the decision process as a result of the additional information. A decision process can be thought of as a simple two-person game. The decision-maker first chooses a particular action, e.g., a particular replacement interval; the other player then chooses a value of a parameter upon which the outcome of the game critically depends, e.g., a failure rate; given the action and the value of the parameter, the decision-maker receives a specific payoff, e.g., a particular in-commission rate. Before playing the game the decision-maker can purchase information about the value of the parameter to be chosen by the other player, e.g., he can, by means of life testing and/or the establishment of a data processing system, purchase information about the unknown failure rate. This information will aid the decision-maker in choosing the preferred action. The appropriate amount of information to purchase depends, of course, on the value of information, in terms of added payoff, relative to its acquisition cost. Usually, acquisition cost is easily estimated. The value of information can be estimated by use of statistical decision theory.

The analysis is based almost exclusively on methods presented in Ref. 9. The methods which are pertinent to the replacement problem are summarized in Sec. 2.2. The application of these methods to the replacement problem is presented in Sec. 2.3. Readers who are familiar with the analysis contained in Ref. 9 should go directly to Sec. 2.3.

2.2 The Value of Information

The application of the following section is based essentially on two concepts of statistical decision theory: the expected value of perfect information and the expected value of sample information. This section is an elaboration of these two key concepts.*

2.2.1 Preliminaries

Let $u(a, \theta)$ represent the utility or satisfaction achieved by a decision-maker when he chooses action a and the state of nature is θ. If the decision-maker knew the state of nature with certainty he would obviously choose the action a which maximized his utility. The state of nature is usually not known with certainty but frequently may be characterized by a random variable $\tilde{\theta}$ with a known probability distribution. It will be assumed that in this circumstance the decision-maker chooses the action which maximizes his expected utility, where the expectation is with respect to the distribution of $\tilde{\theta}$.

The decision-maker is often able to learn more about the random variable $\tilde{\theta}$ before taking a particular action and thereby chooses a better action. This learning may be thought of as an experiment, e,

*For a detailed discussion, see Ref. 11, Chap. 4.

which produces a particular outcome, z. The decision-maker chooses the experiment and nature chooses the outcome. A decision which permits prior experimentation is in fact a two-stage decision problem. The decision-maker chooses an experiment, e, nature determines the outcome of the experiment, z, and the decision-maker receives a payoff, u(e, z), which of course will be negative in most cases. At the second stage the decision-maker chooses an action, a, nature determines the state of nature, θ, and the decision-maker receives a payoff, $u(a, \theta)$. The total payoff is given by

$$u = u(e, z, a, \theta) \tag{2.1}$$

where it is assumed here that the total payoff is decomposible into two parts

$$u = u(e, z) + u(a, \theta). \tag{2.2}$$

With the experiment and action at his disposal, the decision-maker chooses e and a to maximize his utility u.

2.2.2 The Expected Value of Perfect Information

Let a_θ be the act which the decision-maker would choose if θ were the true state of nature, i.e.,

$$u(a_\theta, \theta) \geq u(a, \theta), \text{ all } a. \tag{2.3}$$

Suppose that the state of nature θ is not known with certainty but instead is a random variable $\tilde{\theta}$ with known probability distribution, $p'(\theta)$. Let a' be the act which maximizes the decision-maker's expected utility, i.e., (2.4) $$E'_\theta\, u(a', \tilde{\theta}) \geq E'_\theta\, u(a, \tilde{\theta}) \text{ all } a,$$

where the expectation E'_θ is with respect to the probability distribution $p'(\theta)$.

For a cost of c_∞ the decision-maker can purchase information on the exact value of θ. What is the value of this information? Notice that if the decision-maker acted without this information he would choose a' and anticipate an expected utility of $E'_\theta\, u(a', \tilde{\theta})$ whereas utility of amount $u(a_\theta, \theta)$ will be forthcoming if perfect information is purchased. The value of perfect information for a given state of nature is given by

(2.5) $$v(e_\infty, \theta) = u(a_\theta, \theta) - u(a', \theta),$$

where e_∞ is the experiment which produces perfect information about the state of nature.

This calculation assumes that the state of nature is known. The exact value will of course be available after perfect information is purchased, but it is precisely this purchase which is being questioned. The purchase decision can, however, be based on the expected value of perfect information where this quantity is given by

(2.6) $$v^*(e_\infty) = E'_\theta\, v(e_\infty, \tilde{\theta}) = E'_\theta\, u(\tilde{a}_\theta, \tilde{\theta}) - E'_\theta\, u(a', \tilde{\theta}).$$

The following decision rule may be adopted:

Purchase perfect information if $v^*(e_\infty) \geq c_\infty$

Do not purchase perfect information if $c_\infty > v^*(e_\infty)$.

2.2.3 The Expected Value of Sample Information

Let z denote the outcome of a particular experiment e. Then a_z represents the act which maximizes the decision-maker's expected utility, where the expectation is with respect to the posterior distribution $p''(\tilde{\theta}|z)$, of $\tilde{\theta}$ --the distribution which is in effect after the outcome z is observed. More formally, a_z is defined by the inequality,

$$(2.7) \qquad E''_{\theta|z}\, u(a_z \tilde{\theta}) \geq E''_{\theta|z}\, u(a, \tilde{\theta}), \text{ all a.}$$

The conditional value of the sample information z derived from an experiment e is represented by

$$(2.8) \qquad v(e, z) \equiv E''_{\theta|z}\, U(a_z, \tilde{\theta}) - E''_{\theta|z}\, u(a', \tilde{\theta})$$

where a' is defined as before.

This calculation assumes that the exact value of z is available, i.e., that the experiment has been performed. The decision as to whether or not to experiment cannot be guided by this calculation. However, prior to performing the experiment the expected value v*(e) of the expression v(e, z) can be calculated where the probability distribution over which this expectation is taken is the conditional distribution of the outcome given the experiment, p(z|e). Whether or not the experiment should be performed will depend on the relation between this expectation and the cost c(e) of the experiment:

Perform the experiment if $v^*(e) \geq c(e)$

Do not perform the experiment if $c(e) > v^*(e)$.

2.3 Application to a Replacement Problem

Recent research on the theory of replacement and inspection provides a large class of problems which fit quite naturally within the framework of statistical decision theory.* This section illustrates this relationship by applying statistical decision theory to a simple replacement problem. Similar applications may be executed for the more complex replacement and inspection problems, but since they add little to the understanding of the underlying relationship they are not presented.

2.3.1 A Simple Replacement Problem

The simplest replacement problem is determining an optimal replacement policy for an equipment which cannot be inspected and which fails according to a known failure distribution. A replacement action constitutes a regeneration point, i.e., the future behavior of the equipment is independent of its pre-replacement history. This renewal property implies that the criterion of maximizing time good over an infinite time horizon is equivalent to the criterion of maximizing time good over a single replacement cycle. Assuming that the equipment fails exponentially, the average in-commission rate over a replacement cycle is given by

$$G(N, \lambda) = \frac{(1-e^{-\lambda N})}{\lambda(N+K)}, \tag{3.1}$$

where N is the time between replacements, K is the downtime accompanying the replacement action, and λ is the failure rate. The ratio G may be interpreted as an in-commission rate. Maximizing this expression with

---

*A survey of this work is contained in Ref. 10.

respect to N, the optimal value, N_λ is the solution to the equation

(3.2) $$e^{\lambda N} - \lambda N = 1 + \lambda K.$$

Applying a quadratic approximation to $e^{\lambda N}$ the equation reduces to

(3.3) $$N_\lambda = \sqrt{2K/\lambda}\ .$$

In many applications the parameter λ is not known with certainty but instead is a random variable $\tilde{\lambda}$ with a known probability distribution. Using Bayes Theorem to revise this distribution as sampling information accumulates, a natural adaptive criterion is to choose the replacement interval N' which satisfies the expression

(3.4) $$\underset{N}{\text{Max}} \int_0^\infty G(N, \lambda)\ \phi'(\tilde{\lambda})d\lambda$$

where $\phi'(\tilde{\lambda})$ is the probability density function of the distribution which summarizes all available information on $\tilde{\lambda}$. Assuming that the random variable $\tilde{\lambda}$ possesses a Gamma distribution, it has been shown (Ref. 5) that the adaptive replacement interval is the solution to the equation

(3.5) $$\left(1 + \frac{N}{t'}\right)^{r'} - r't'^{r'-1} N - t'^{r'-1} K(r'-1) = 0$$

where r' and t' are the parameters of the relevant Gamma distribution. If the decision-maker has no control over the generation of information, or equivalently, if this information is free, the decision-maker simply revises $\phi'(\lambda)$ as new information accrues and establishes a new replacement interval satisfying Eq. (3.4).

Very often the decision-maker has some control over the generation of information and, furthermore, this information is almost never free. When this is true, a more complicated decision rule is required. Given that a particular sample had been taken, the decision-maker would choose the adaptive replacement interval according to Eq. (3.4). However, now the decision-maker must also decide whether or not to collect more information and if so, how much more. Two situations will be distinguished. In the first, the decision-maker has two alternatives: (1) He may, for a fixed cost, purchase perfect information, obtaining the exact value of the unknown parameter. His future replacement decisions will thereby be fully informed; or (2) he may purchase no additional information and base his future replacement decisions on currently available information. These alternatives are frequently approximated in practice when the decision-maker chooses whether or not to purchase data processing equipment. The use of this equipment may produce almost perfect information. The purchase decision depends on the cost of the equipment relative to the expected value of the more informative decisions made possible by the purchase. The second situation again involves a choice between two information levels. The decision-maker may base his future replacement decisions on information already available or he may take a sample of a specified size and modify his future replacement decisions accordingly. The cost of information is zero or some fixed amount depending upon whether the first or second alternative is selected. This choice is often encountered in practice. For example, suppose that the equipment being replaced is a ballistic missile. All of the parameters affecting the replacement policy are known with certainty with the exception of a single critical failure rate, λ.* The failure rate is a random variable with a known probability distribution. The decision-maker has two alternatives at his disposal. For a known cost he may establist a test program which will produce information about the failure rate, λ. The replacement policy would then be conditioned by the results of this test. Alternatively, the decision-maker may base the replacement policy on whatever information is freely available. As before, the choice between these alternatives depends on the cost of the test program in comparison with the expected value--measured in missile effectiveness--of the additional information.

2.3.2 The Expected Value of Perfect Information

All variables affecting the replacement decision are known with certainty except the failure rate $\tilde{\lambda}$ which is a random variable with a Gamma density function

$$\phi'(\tilde{\lambda}) = \begin{cases} \dfrac{\lambda^{r'-1} t'^{r'} e^{-t'\lambda}}{\Gamma(r')} & , \lambda \geq 0 \\ 0 & , \text{elsewhere} \end{cases} \tag{3.6}$$

where r' and t' are known parameters. This density function incorporates all information accessible to the decision-maker. Using notation analogous to that of Sec. 2.2, let N' denote the replacement interval which maximizes the expected in-commission rate, that is,

$$E_\lambda' \, G(N', \tilde{\lambda}) \geq E_\lambda' \, G(N, \tilde{\lambda}), \text{ all } N \tag{3.7}$$

---

*See Ref. 4.

where the expectation E'_λ is taken over the distribution characterized by $\phi'(\tilde{\lambda})$. More specifically, the replacement interval N' is the value of N satisfying the expression (3.4).

If the decision-maker purchases no additional information the replacement interval would be set equal to N' and the expected in-commission rate would be $E'_\lambda\ G(N', \tilde{\lambda})$. Alternatively, the decision-maker has the option of purchasing perfect information on the failure rate λ at a cost of c_∞. If this information were acquired, the decision-maker would be in a position to choose the replacement interval which satisfied Eq. (3.1)--replacement interval would be set equal to N_λ and the in-commission rate would be $G(N_\lambda, \lambda)$. For a given λ, the value of perfect information can be expressed as

$$(3.8) \qquad v(e_\infty, \lambda) = G(N_\lambda, \lambda) - G(N', \lambda)$$

where e_∞ signifies the experiment which yields perfect information about the parameter λ.* The value of perfect information when the value of the failure rate is λ is simply the difference between the in-commission rate when the replacement interval is N_λ and the highest expected in-commission rate achievable with available information. This calculation assumes that the failure rate is known. This, of course, will be the case if perfect information is purchased. Since it is precisely this purchase which is being evaluated, the purchase decision cannot be founded on the value of perfect information but may be founded on the expected value of perfect information where this latter quantity is given by

$$(3.9) \qquad v^*(e_\infty) = E'_\lambda\ G(\tilde{N}_\lambda, \tilde{\lambda}) - E'_\lambda\ G(N', \tilde{\lambda}).$$

---

*The comparison of the value of perfect information with its cost presupposes that $G(N, \tilde{\lambda})$, the in-commission rate, has been converted to its dollar equivalent. This is easily accomplished since maximizing $G(N, \lambda)$ is

Using Eq. (3.1) and Eq. (3.3), this quantity can be approximated by

$$(3.10) \quad v^*(e_\infty) \cong \int_0^\infty \frac{1-e^{-\sqrt{2K\lambda}}}{\sqrt{2K\lambda}+K\lambda}\, \phi'(\tilde{\lambda})d\lambda - E'_{\tilde{\lambda}}\, G(N', \tilde{\lambda})$$

which is easily computed.

In principle, the expected value of perfect information should be calculated before any additional information is purchased. If the cost of information which is less than perfect exceeds the value of perfect information, then clearly the information should not be purchased.* Of course, it may be disadvantageous to purchase imperfect information even if the value of perfect information exceeds the cost of the imperfect information. As indicated previously, the decision-maker is often able to purchase almost perfect information by the installation of data processing equipment. If the sole purpose of this installation is to facilitate the replacement decision, then the appropriate rule is:

Install the data processing equipment if $v^*(e_\infty) \geq c_\infty$.

Do not install the data processing equipment if $v^*(e_\infty) < c_\infty$

where c_∞ is the acquisition cost.

2.3.3 The Expected Value of Sampling Information to the Replacement Decision

The decision-maker is oftentimes unable to perform a "perfect information" experiment. In the replacement problem the alternatives

---

equivalent to minimizing the cost per unit of in-commission time per unit time. (See Ref. 5.) In the sequel, comparative statements between the cost and value of information assume that this conversion has been carried out.

*Needless to say, this assumes that information of value to more than one decision process is evaluated accordingly. For example, failure rate information may be of value to a supply (stockage) decision as well as to the maintenance (replacement) decision.

are often: (1) establish a replacement interval based on available information by executing a sampling experiment which terminates when r failures are observed--a Gamma sampling scheme.* The cost of the sampling information is denoted by c_z, where z signifies the information revealed by the experiment. Essentially the same analysis as before can be employed to appraise each alternative.

Let N_z represent the replacement interval which maximizes the expected in-commission rate where the expectation is with respect to the posterior distribution of the failure rate--the distribution in effect after the experiment has been executed. The posterior distribution is characterized by a Gamma density function, $\phi''(\tilde{\lambda}|z)$ with parameters.

$$r'' = r' + r$$

$$t'' = t' + t .$$

More precisely, N_z is defined by the inequality

$$E''_{\lambda|z}\, G(N_z, \tilde{\lambda}) \geq E''_{\lambda|z}\, G(N, \tilde{\lambda}), \text{ all } N \tag{3.11}$$

where the double prime indicates that the expectation is over the posterior distribution of λ. The conditional value of the sample information is given by

$$v(e, r) \equiv E''_{\lambda|z}\, G(N_z, \tilde{\lambda}) - E''_{\lambda|z}\, G(N', \tilde{\lambda}) \tag{3.12}$$

where N' is defined as before. The calculations are greatly facilitated if the function G(N, λ) is linear in λ. A quadratic approximation to $e^{-\lambda N}$ in Eq. (3.1) produces the desired linearity,

*See Ref. 11, p. 279.

(3.13) $$G(N, \lambda) \cong \frac{2N-N^2\lambda}{2(N+K)} .$$

This linearity implies the equation

(3.14) $$E_{\lambda}\, G(N, \tilde{\lambda}) = G(N, \bar{\lambda})$$

where $\bar{\lambda}$ is the mean of the distribution. The expected value of the in-commission rate over the distribution of λ is exactly equal to the in-commission rate with the variable λ replaced by its expected value, $\bar{\lambda}$. The replacement interval N' is now defined by the equation

(3.15) $$G(N', \bar{\lambda}') = \max_{N} G(N, \bar{\lambda}')$$

where $\bar{\lambda}'$ is the mean of the prior distribution $\phi'(\lambda)$. Similarly, N_z is defined by the equation

(3.16) $$G(N_z, \bar{\lambda}'') = \max_{N} G(N, \bar{\lambda}'')$$

where $\bar{\lambda}''$ is the mean of the posterior distribution $\phi''(\lambda|z)$. The conditional value of the sample information is now definable as

(3.17) $$v(e, \bar{\lambda}'') = \max_{N} G(N, \bar{\lambda}'') - G(N', \bar{\lambda}'').$$

This calculation assumes that the expected value of the posterior distribution $\bar{\lambda}''$ is known, whereas in fact it will not be known until after the experiment has been performed. However, prior to the experiment it is possible to calculate the expected value of the sample information where the expectation is with respect to the prior distribution of the posterior mean, $\bar{\lambda}''$--the "preposterous" distribution.* The expected value of the sample information is given by

$$v^*(e) = E_{\bar{\lambda}''|e}\, v(e, \tilde{\bar{\lambda}}'') \tag{3.18}$$

where $\tilde{\bar{\lambda}}''$ indicates that prior to the experiment the posterior mean $\bar{\lambda}''$ is a random variable. It is shown in Ref. 11 that when the sampling is Gamma the preposterous distribution of $\bar{\lambda}''$ is beta with F_β denoting the beta cumulative distribution function. The mean and variance of this distribution are, respectively,

$$E(\bar{\lambda}'') = r'/t'$$

$$\mathrm{Var}\ (\bar{\lambda}'') = \frac{r^2}{(r+r'+1)t'^2}\ .$$

Therefore, utilizing a quadratic approximation Eq. (3.13), Eq. (3.18) can be written

$$v(e) \cong \int_0^\infty G_a(N_z, \tilde{\bar{\lambda}}'')dF_\beta - \int_0^\infty G_a(N', \tilde{\bar{\lambda}}'')dF_\beta \tag{3.19}$$

where the subscript a indicates that the in-commission rates are being linearly approximated.

---

*See Ref. 11, p. 104.

3. THE COST AND VALUE OF INFORMATION IN INVESTMENT PROBLEMS

3.1 Introduction

In this section the methodology of optimal stopping[1,6,8] and renewal programming[3] is used to analyze the informational aspects of several investment problems.[13] Two different types of investment decisions are considered. In the first, investment opportunities are assumed to arise according to a known probability mechanism. A fixed amount of funds are available for investing and the decision-maker must decide between investing in a known opportunity and waiting for the next uncertain opportunity. Investment opportunities are of the "now or never" variety in that a rejected opportunity can never be reconsidered. Many practical problems possess this "now or never" feature. The second kind of investment decision differs from the first in that the decision-maker stops the investment process by investing in the best of the observed opportunities. This investment climate is more generous than the first and frequently occurs in practice.

The general stochastic investment process is governed by three distinct probability laws. In addition to analyzing the general process, three specializations are also considered. The description of the general process and each of the specializations is facilitated by referring to Table 1. The basic probability mechanism is a discrete parameter Markov chain with transition matrix $|P_{ij}|$, $i,j = 1,2,\ldots,k$. The movement from investment opportunity i to investment opportunity j is controlled by the Markov chain. The movement from i to j is, however, neither instantaneous nor of fixed duration, but is a random variable, τ_{ij}, with delay distribution, F_{ij}. Imbedding this delay distribution in the Markov chain gives rise to a semi-Markov process (second row, first column of Table 1 below), so called because the

Table 1

STOCHASTIC PROCESSES IMBEDDED IN A DISCRETE PARAMETER MARKOV CHAIN

		Delay Distribution, F_{ij}: Non-Degenerate	Delay Distribution, F_{ij}: Degenerate
Opportunity Distribution, G_i	Non-Degenerate	Generalized Semi-Markov Process	Karlin Process
	Degenerate	Semi-Markov Process	Simple Markov Process

process is Markovian only at the instants when states are selected. Finally, the investment opportunity associated with a particular state i is also a random variable with probability distribution, G_i. When this distribution is coupled with the Markov chain, a quasi semi-Markov process ensues. Karlin was the first to introduce this process and in the sequel it is called the Karlin process (first row, second column of Table 1 above). In this paper, investment opportunities are described by a single number, the rate of return. Accordingly, the random variable, ρ_i, associated with the probability distribution, G_i, is a rate of return.

The first stochastic model discussed is one in which there is only one state with a rate of return distribution, G,--the Markov chain is degenerate and concentrated on a single state. Furthermore, F, the distribution of the time-between-investment opportunities, is also degenerate, all of its mass being concentrated on a single value. This probability model is applied to each type of investment decision and optimal investment policies are derived for both fixed and infinite planning horizons.

The second stochastic model is the semi-Markov process in which the Markov chain and delay distribution are full-blown but the rate of return distribution is degenerate, all of its mass concentrated on a single value, and $G_i = G_j$ all i and j.

The third stochastic model is the Karlin process. In this process the Markov chain and rate of return distribution are full-blown, but the delay distributions are equal and degenerate.

The final model is the general semi-Markov process in which none of the probability laws is degenerate. Notice that if both the delay distributions and the rate of return distributions are degenerate, the resulting probability process has a simple Markovian structure.

In the analysis the decision-maker is required to pay a fixed fee for each drawing from the rate of return distribution. This informational cost is incurred even when the rate of return distribution is degenerate. Discounting considerations give rise to a second kind of informational cost; other things equal, the decision-maker prefers to invest as soon as possible because of the discounting factor. The criterion used to evaluate alternative investment strategies is the net present value of the investment, that is, the present value of the investment less the cost of information. The optimal investment policy maximizes the present value.

Before turning to the analysis it might be useful to exemplify the general semi-Markov process. The decision-maker is a recent Ph.D. in economics and is in the midst of selecting a job, i.e., investing his human capital. Three distinct groups are bidding for his talents, government, academic and business. The offer made by each group is a random variable with probability distribution, G_i, $i = 1,2,3$, where 1,2, and 3 denote government, academic and business, respectively. Given that an offer has just been made by group i, the probability of an offer from group j is given by the Markov transition matrix $|P_{ij}|$, $i,j = 1,2,3$. Finally, the waiting time between an offer from group i and an offer from group j is also a random variable with probability distribution, F_{ij}, $i,j = 1,2,3$. No income is earned until an offer is accepted and this constitutes the cost of information. Assuming that offers have been standardized to reflect preferences, e.g., \$12,000 from business may be equivalent to \$10,000 from government which is equivalent to \$9,000 from academia, the Ph.D. desires to accept that offer that maximizes his standardized income. Two acceptance regimes are permissible. The offers may be accepted or rejected as they occur, with an acceptance terminating the process; or all acquired offers may be outstanding until an acceptance occurs when all other offers are rejected and the process ends.

3.2 Investments are Accepted or Rejected as They Occur

The first investment problem to be discussed is the case where investment opportunities occur randomly and immediately must be accepted or rejected by the decision-maker. An investment opportunity is completely characterized by a single rate of return, ρ. The objective of the decision-maker is to maximize the discounted value of the funds that are available for investment. This problem will be analyzed for four different probability processes: investment opportunities may be generated by a single stationary probability distribution, a semi-Markov process, a Karlin process, or a generalized semi-Markov process.

3.2.1 Investment Opportunities are Periodically Drawn from a Stationary Probability Distribution

An investment opportunity is a random variable with a known stationary probability distribution. Each period the decision-maker pays a fixed fee, c, particular investment opportunity is presented and is either accepted or rejected. Let K and ρ denote respectively the amount of funds available for investment and the rate of return associated with an investment opportunity. The random variable ρ has probability distribution, G.

3.2.1.1 Optimal Investment Policy for a Finite Number of Periods*

First assume that the decision-maker must accept one of the first N investment opportunities. Let K_N denote the maximum discounted expected value when one of the N opportunities is selected according to an optimal policy. A policy, $\Psi(\rho,N)$, is the probability of accepting the first of the N opportunities when ρ is its rate of return. If the first investment opportunity is accepted or rejected

*The following argument is essentially the same as that of Karlin, Ref. 6, p. 149. Also see Ref. 2.

in accordance with $\Psi(\rho, N)$ and from then on an optimal policy is followed, the expected return is

$$H(\Psi, N-1) = \int_0^\infty \left[\text{Max}\left(\frac{\rho K}{\rho_0} - c,\ K-c \right) \right] \Psi(\rho, N)\, dG(\rho)$$

(3.1)

$$+ \alpha R_{N-1} \int_0^\infty (1-\Psi(\rho, N))\, dG(\rho)$$

$$= \int_0^\infty \left[M - \alpha R_{N-1} \right] \Psi(\rho, N)\, dG(\rho) + \alpha R_{N-1}$$

where $M = \text{Max}\left(\frac{\rho K}{\rho_0} - c,\ K-c \right)$, $N \geq 1$, $R_0 = 0$, ρ_0 is the opportunity rate of return--the return always available outside this process, and $\alpha = \frac{1}{1+\rho_0}$ is the discount factor. By definition, maximizing $H(\Psi, N-1)$ with respect to the policy parameter Ψ yields R_N, the discounted income stream obtainable from following an optimal stopping policy when N investment opportunities are available:

$$\underset{\Psi}{\text{Max}}\ H(\Psi, N-1) = R_N .$$

It follows that the optimal investment policy is given by

(3.2)
$$\Psi^*(\rho, N) = \begin{cases} 1 \text{ (accept)},\ M > \alpha R_{N-1} \\ 0 \text{ (continue searching)},\ \alpha R_{N-1} > M. \end{cases}$$

The optimal policy is degenerate in the sense that the probability of accepting a particular opportunity is either zero or one. Calculation of R_N is achieved by substituting Eq. (3.2) into Eq. (3.1) and proceeding recursively:

$$(3.3) \qquad R_N = \int_{\frac{\alpha \rho_0 R_{N-1}}{K}}^{\infty} \left[M - \alpha R_{N-1} \right] dF(\rho) + \alpha R_{N-1} .$$

The expected value of additional information is simply:

$$(3.4) \qquad EVI = \alpha R_{N-1} - M.$$

3.2.1.2 Optimal Investment Policy for an Infinite Number of Periods

Letting R denote the discounted expected value when an optimal policy is followed,

$$(3.5) \qquad R = \sup_{0 \le \Psi(\rho) \le 1} \left\{ \int_0^{\infty} \Psi \, M \, dF(\rho) + \alpha R \int_0^{\infty} (1-\Psi) \, dF(\rho) \right\} .$$

It follows that the optimal policy is given by

$$(3.6) \qquad \Psi^* = \begin{cases} 1 \text{ (accept)}, & M > \alpha R \\ 0 \text{ (reject)}, & \alpha R > M \end{cases} .$$

Using Eq. (3.6) in Eq. (3.5) yields the relation

$$(3.7) \qquad R = \int_{\frac{\alpha \rho_0 R}{K}}^{\infty} \Psi(M - \alpha R) \, dF(\rho) + \alpha R$$

which possesses a unique root.[6]

3.2.2 Investment Opportunities Are Generated by a Semi-Markov Process

The same "now or never" investment problem is considered when investment opportunities arise stochastically according to a semi-Markov process. It will be recalled that a semi-Markov process is characterized by a Markov chain together with a delay distribution. Each state of the process contains a non-random investment opportunity, ρ. The decision to reject a particular opportunity is accompanied by simultaneous drawings from the delay distribution and the Markov chain that give, respectively, the wait until the next opportunity and its value. The cost of these drawings is again denoted by c.

3.2.2.1 Optimal Investment Policy for a Finite Number of Periods

The decision-maker is required to accept one of the first N investment opportunities, that is, there are at most N transitions. Let $R_i(N, \alpha)$ denote the expected discounted value from following an optimal investment policy when the process begins in state i with rate of return ρ_i and the discounting factor is α. The time between transitions, τ_{ij}, is a random variable with known probability distribution, F_{ij}. As before, K funds are available for investment and investments are accepted or rejected. If investment opportunity ρ_i is accepted the process is terminated and the decision-maker obtains an income stream with capital value, $M = \text{Max}\left(\frac{\rho_i K}{\rho_0} - c, K-c\right)$, where ρ_0 is the opportunity rate of return, i.e., the rate at which K could be invested before the process began.

Alternatively, if ρ_i is rejected, the expected discounted return from following an optimal policy thereafter may be expressed as:

$$(3.8) \qquad \Pi_i = \sum_{j=1}^{k} p_{ij} \int_0^\infty e^{-\alpha t}\, dF_{ij}(t)\, R_j(N-1, \alpha).$$

Hence,

$$R_i(N, \alpha) = \text{Max} \begin{cases} \text{accept}, & M > \Pi_i \\ \text{reject}, & \Pi_i > M \end{cases} . \tag{3.9}$$

The expected value of additional information is given by

$$EVI = \Pi_i - M \tag{3.10}$$

and the process is continued or terminated if EVI is positive or negative, respectively.

3.2.2.2 Optimal Investment Policy for an Infinite Number of Periods

An infinite horizon investment process places no restrictions on the number of transitions. The decision-maker can wait as long as he wishes before accepting an investment opportunity. Let $R_i(\alpha)$ denote the maximum expected discounted value when the investment process begins in state i and the discount factor is α. Then using the same notation as before,

$$\begin{aligned} R_i(\alpha) &= \sup_{0 \le \Psi \le 1} \left\{ M \, \Psi_i(\rho_i, \alpha) + \sum_{j=1}^{k} p_{ij} \, R_j \int_0^\infty e^{-\alpha t} \, dF(t) \left(1 - \Psi_i(\rho_i, \alpha)\right) \right\} \\ &= \sup_{0 \le \Psi \le 1} \left\{ (M - \Pi_i) \, \Psi_i(\rho_i, \alpha) + \Pi_i \right\} \end{aligned} \tag{3.11}$$

where

$$\Pi_i = \sum_{j=1}^{k} p_{ij} R_j \int_0^\infty e^{-\alpha t} dF(t).$$

It follows that the optimal policy is characterized by the k-triple,

$$(\Psi_1^*, \Psi_2^*, \ldots, \Psi_k^*)$$

where

$$(3.12) \qquad \Psi_i^* = \begin{cases} 1 \text{ (accept investment i)}, & M > \Pi_i \\ 0 \text{ (reject investment i)}, & \Pi_i > M. \end{cases}$$

Substituting Eq. (3.12) into Eq. (3.11) yields the system of equations

$$(3.13) \qquad R_i = \text{Max}(M, \Pi_i)$$

that can be used to calculate the k critical numbers, Π_i, $i = 1,2,\ldots,k$.

The optimal policy has the same structure as before. The decision-maker simple compares the known discounted value of investment opportunity i with the expected discounted value of pursuing an optimal policy from then on. If the difference is negative, opportunity i is accepted; otherwise, the process continues. The expected value of additional information is given by

$$(3.14) \qquad EVI = (\Pi_i - M).$$

3.2.3 Investment Opportunities are Generated by a Karlin Process

In this process the movement from state to state is again governed by a Markov chain. The transition from state to state is of fixed duration. As before, an investment opportunity is associated with each state. However, each of these opportunities is a random variable with a known probability distribution, F_i, $i = 1,2,\ldots,k$. The process is similar to that of Sec. 3.2.1 except that the probability distribution of investment opportunities need not be stationary, that is, F_i and F_j may differ. This probability process has been used to analyze inventory problems possessing known non-stationary demand distributions.[7] The same kind of non-stationary environment is undoubtedly present in many practical investment problems.

3.2.3.1 Optimal Investment Policy for a Finite Number of Periods

The decision-maker is first required to accept one-of-the first N investment opportunities. Let $R_i(N, \alpha)$ denote the expected discounted return from following an optimal policy when the process begins in state i. If the process is terminated by the acceptance of investment opportunity ρ_i, a discounted value of $M = \text{Max}\left(\frac{\rho_i K}{\rho_0} - c, K - c\right)$ is realized. Alternatively, continuation of the process according to an optimal policy yields an expected discounted value that may be expressed as

$$\Pi_i = \alpha \sum_{j=1}^{k} p_{ij} R_j(N-1, \alpha). \tag{3.15}$$

The optimal policy for the i^{th} state when N opportunities remain is characterized by the following decision rule:

$$\begin{aligned} &\text{accept } \rho_i \text{ if } M > \Pi_i \\ &\text{reject } \rho_i \text{ if } \Pi_i > M \end{aligned} \tag{3.16}$$

where Π_i can be calculated recursively from the relation

$$(3.17) \qquad R_i(N, \alpha) = \text{Max} \{M, \Pi_i\}, \; i = 1,2,\ldots k.$$

This calculation is facilitated by the presence of only two decisions at each point.(3) The expected value of additional information is simply

$$EVI = \Pi_i - M .$$

3.2.3.2 Optimal Investment Policy for an Infinite Number of Periods

The case where investment opportunities are unlimited is now considered. The analysis closely parallels that of Sec. 3.2.2.2.* A policy is described by a k-triple $(\Psi_1, \Psi_2, \ldots, \Psi_k)$, where Ψ_i is the probability of accepting the i^{th} investment opportunity. The expected discounted value of following an optimal policy when the process begins in state i is denoted by R_i. Then

$$(3.18) \qquad R_i = \sup_{0 \le \Psi_i \le 1} \left\{ \int_0^\infty M\Psi_i(\rho) + \alpha \sum P_{ij} R_j \int_0^\infty (1-\Psi_i) \, dF_i(\rho), \; i = 1,2,\ldots,k. \right\}$$

Letting $\Pi_i = \alpha \sum P_{ij} R_j$, the optimal policy is given by

$$(3.19) \qquad \Psi_i^* = \begin{cases} 1 \text{ (acceptance)}; & M > \Pi_i \\ 0 \text{ (rejection)}; & M < \Pi_i \end{cases} \qquad i = 1,2,\ldots,k.$$

*For a complete discussion of the analysis, see Karlin.(6)

Calculation of the k critical numbers, Π_i, i = 1,2,...,k, can be accomplished by solving the set of equations

$$(3.20) \qquad R_i = \Pi_i + \int_{\frac{\Pi_i \rho_0}{K}}^{\infty} (M - \Pi_i)\, dF_i(\rho), \; i = 1,2,\ldots,k.$$

The expected value of additional information is given by

$$(3.21) \qquad EVI = \Pi_i - M.$$

3.2.4 The Generalized Semi-Markov Process

The final and most general stochastic process is described by a Markov chain and two probability distributions, a delay distribution and a distribution of investment opportunities. At each decision point the decision-maker can accept the known rate of return, ρ_i, associated with state i, in which case the process is stopped. If, however, the decision-maker rejects the known opportunity then the succeeding state j is chosen according to the transition matrix $|P_{ij}|$. This is followed by a draw from the delay distribution, F_{ij}, that determines the waiting time until state j is reached. Finally, when state j is achieved, a sample of one is drawn from G_j, the distribution of investment opportunities, and ρ_j, the resulting rate of return is either accepted or rejected. The three previously described probability processes are special cases of the generalized semi-Markov process.

3.2.4.1 Optimal Investment Policy for a Finite Number of Periods

The decision-maker is required to accept one of the first N investment opportunities. The discounted value of accepting investment opportunity i is given by M and the expected discounted value of continuing the process in optimal fashion is

$$(3.22)\quad \Pi_i = \sum_{j=1}^{k} p_{ij} \int_0^\infty e^{-\alpha t}\, dF_{ij}(t)\, R_j(N-1, \alpha),\ i = 1,2,\ldots,k,$$

where $R_i(N, \alpha)$ can be calculated recursively from the relation

$$(3.23)\quad R_i(N, \alpha) = R_i(N-1, \alpha) + \int_{\frac{\Pi_i \rho_0}{K}}^{\infty} (M - \Pi_i)\, dF_i(\rho),\ i = 1,2,\ldots,k.$$

The optimal investment policy for the $i^{\underline{th}}$ state when N opportunities remain is given by

$$(3.24)\quad \Psi_i^*(N, \alpha) = \begin{cases} 1 \text{ (acceptance)}; & M > \Pi_i \\ 0 \text{ (rejection)}; & \Pi_i > M. \end{cases}$$

The expected value of additional information is denoted by

$$(3.25)\quad EVI = \Pi_i - M.$$

3.2.4.2 Optimal Investment Policy for an Infinite Number of Periods

The analysis is identical to that of Sec. 3.2.3.2 with the exception that Π_i is now expressed as

$$(3.26)\quad \Pi_i = \sum_{j=1}^{k} p_{ij} \int_0^\infty e^{-\alpha t}\, dF_{ij}(t)\, R_j .$$

3.3 Optimal Investment Policy When the Best Observed Opportunity May be Selected

In many practical situations the decision-maker is permitted to accumulate investment opportunities and terminates the investment process by selecting the best observed opportunity. The parameters affecting this decision process are identical to those of the "now or never" investment problem and the same definitions are employed. Two different probability processes are analyzed. First, investment opportunities occur periodically according to a known probability distribution. Secondly, investment opportunities are generated by a generalized semi-Markov process.

3.3.1 Investment Opportunities are Periodically Drawn from a Stationary Probability Distribution

3.3.1.1 Optimal Investment Policy for a Finite Number of Periods

Let $R(N, \rho_m)$ denote the discounted expected value when one of the remaining N opportunities is optimally selected and ρ_m is the maximum rate of return so far observed. Then

$$(3.27) \quad R(N, \rho_m) = \text{Max} \left\{ \left(\frac{\rho_m K}{\rho_0} - c \right), \left(1-G(\rho_m) \right) \left(R(N-1, \rho_m) \right) + \int_{\rho_m}^{\infty} R(N-1, \rho) \, dG(\rho) \right\}$$

and the optimal policy is given by

$$(3.28) \quad \Psi^*(N, \rho_m) = \begin{cases} 1 \text{ (acceptance)}, & M_0 > \Pi \\ 0 \text{ (rejection)}, & \Pi > M_0 \end{cases}$$

where

$$M_0 = \frac{\rho_m K}{\rho_0} - c$$

and

$$\Pi = \left(1-G(\rho_m)\right)\left(R(N-1, \rho_m)\right) + \int_{\rho_m}^{\infty} R(N-1, \rho)\, dG(\rho).$$

The expected value of additional information is

(3.29) $$EVI = \Pi - M_0 .$$

3.3.1.2 Optimal Investment Policy for an Infinite Number of Periods

Let $R(\rho_m)$ denote the discounted expected value when the maximum observed rate of return is ρ_m and the decision-maker pursues an optimal investment policy. Termination of the process guarantees a discounted value of

(3.30) $$M_0 = \frac{\rho_m K}{\rho_0} - c .$$

On the other hand, if the process is continued, the expected return from proceeding optimally is given by

(3.31) $$\Pi = R(\rho_m)\, G(\rho_m) + \int_{\rho_m}^{\infty} R(\rho)\, dG(\rho).$$

Clearly,

$$R(\rho_m) = \text{Max}\ (M_0, \Pi) .$$

MacQueen and Miller[8] have shown that the unique solution to this equation is given by

$$R(\rho_m) = \text{Max}\ (\rho_m,\ \rho^*) \tag{3.33}$$

where

$$\rho^* = \frac{\int_{\rho^*}^{\infty} \rho\ dG(\rho) - c}{\int_{\rho^*}^{\infty} dG(\rho)} .$$

3.3.2 The Generalized Semi-Markov Process

3.3.2.1 Optimal Investment Policy for a Finite Number of Periods

The argument is similar to that of Sec. 3.3.1.1 except that

$$R_i(N,\ \rho_m) = \text{Max}\left(\frac{\rho_m K}{\rho_0} - c,\ \Pi_i\right)$$

and

$$\Pi_i = \sum_{j=1}^{k} p_{ij} \int_0^{\infty} e^{-\alpha t}\ dF_{ij}(t) \left[R_j(N-1,\ \rho_m)\ G_j(\rho_m) + \int_{\rho_m}^{\infty} R_j(N-1,\ \rho)\ dG_j(\rho) \right] ,\ i = 1,2,\ldots,k. \tag{3.34}$$

4. CONCLUSIONS

This paper has presented methods for measuring the economic value of information for two different stopping environments. The decision process occurring in the first environment is repetitive, while the associated probability mechanism is unknown. An equipment replacement problem was used to exemplify this class of stopping problems. Some recent results of statistical decision theory were used to assess the merits of collecting more information about the equipment's failure distribution. This evaluation was performed only for a simple replacement problem, but similar methods can be applied to more complicated scheduling problems. Optimal scheduling policies have been obtained for many problems in which the failure mechanism is known. The application presented here assumes that this optimal structure is preserved when the probability mechanism is not known.

Decision processes comprising the second class of stopping problems possess a terminating structure. The decision-maker is assumed to know the probability law that generates investment opportunities. However, each successive investment opportunity entails informational costs. Given the known opportunities, it may be preferable to select one of them rather than incurring further costs. The methods of dynamic programming were used to determine when investment should occur. Two different investment processes were considered. In the first, investment opportunities were accepted or rejected as they occurred; accumulation of investment opportunities was not permitted. In the second investment process, the decision-maker remembers the best of the revealed opportunities and must choose between accepting it or waiting and purchasing another piece of information.

REFERENCES

1. Breiman, Leo, Stopping Rule Problems, Working Paper No. 7, Western Management Science Institute, University of California, Los Angeles, May 31, 1962.

2. Fisher, James L., "A Class of Stochastic Investment Problems," Operations Research, Vol. 9, No. 1, January-February 1961, pp. 53-65.

3. Jewell, William S., "Markov Renewal Programming I and II," Operations Research, Vol. 11, No. 6, November-December 1963, pp. 938-971.

4. Jorgenson, D. W., and J. J. McCall, "Optimal Replacement Policies for a Ballistic Missile," Management Science, Vol. 9, No. 3, April 1963.

5. -----, "Optimal Scheduling of Replacement and Inspection," Operations Research, October 1963.

6. Karlin, Samuel, "Stochastic Models and Optimal Policy for Selling an Asset," Studies in Applied Probability and Management Science, K. J. Arrow, S. Karlin and H. Scarf, editors, Stanford University Press, Stanford, California, 1962.

7. -----, "Optimal Policy for Dynamic Inventory Process with Stochastic Demands Subject to Seasonal Variations," J. Soc. Indust. Appl. Math., Vol. 8, 1960, pp. 611-629.

8. MacQueen, J., and R. G. Miller, "Optimal Persistence Policies," Operations Research, Vol. 8, No. 3, 1960, pp. 362-380.

9. Marschak, J., "Problems in Information Economics: General Concepts and a Case of Increasing Returns to Information," in Basic Research in Management Controls, edited by Bonini, Jaedicke, and Wagner, Proceedings of Conference at Stanford Graduate Business School, 1963.

10. McCall, J. J., "Optimal Maintenance Policies for Stochastically Failing Equipment: A Survey," Report 6330 (ICMS No. 33) of The International Center for Management Science of The Netherlands School of Economics, Rotterdam, 1963.

11. Raiffa, Howard, and Robert Schlaifer, Applied Statistical Decision Theory, Graduate School of Business Administration, Harvard University, Boston, 1961.

12. Smith, W. L., "Renewal Theory and Its Ramifications," Journal of the Royal Statistical Society, Series B, Vol. 20, No. 2, pp. 243-302.

13. Solomon, Ezra, editor, The Management of Corporate Capital, The Free Press, Glencoe, Illinois, 1959.

14. Stigler, George, "The Economics of Information," Journal of Political Economy, Vol. LXIX, June 1961, pp. 213-225.