

FUNDING CRITERIA FOR RESEARCH, DEVELOPMENT, AND EXPLORATION PROJECTS

BY KEVIN ROBERTS AND MARTIN L. WEITZMAN¹

The sequential nature of activities like research, development, or exploration requires optimal funding criteria to take account of the fact that subsequent funding decisions will be made throughout the future. Thus, there is a continual possibility of reviewing a project's status, based on the latest information. After setting up a model to capture this feature, optimal funding criteria are investigated. In an important special case, an explicit formula is derived. As well as throwing light upon the nature of development activities, the analysis is also relevant to the general theory of information gathering processes.

1. INTRODUCTION

CERTAIN ECONOMIC ACTIVITIES have such special features that the usual "investment criteria" are rendered practically irrelevant. Examples are research, development, and exploration. A real need exists to have simple operational funding rules for such processes which take account of their sequential nature. The present paper is an attempt to characterize when it is optimal to fund a sequential project.

Our point of departure is a single research, development, or exploration project which, by the usual partial equilibrium assumptions, can be analyzed in isolation from the rest of the economy. In the course of undertaking such a project, more and more information is continually being revealed about potential benefits. As it is generally possible to make decisions like backing out if prospects appear unfavorable in a pay-as-you-go project, a rational decision maker will wish to systematically exploit accumulated information.

Throughout this paper, it will be convenient to work with a mathematical abstraction, idealization, or model which we call a *Sequential Development Project* (SDP). This concept is meant to embody some essential features of research, development, or exploration processes: costs are additive; benefits are received only at the termination of the project; there is always the possibility of discontinuing the project altogether.

A basic assumption which will be made is that we can describe where we are in a SDP by an index called its stage or step. As the SDP moves through its various stages, more costs must be paid out, but greater information is amassed because more outcomes of particular stages have been realized.

Benefits in a SDP are unknown or uncertain at each step and are received only at termination. As the SDP proceeds, more information about benefits is accumulated with each step. Mathematically, the distribution of final benefits shifts as results turn out better or worse than anticipated and, at the same time, the

¹We would like to thank Michael Rothschild for his useful comments and John Norbury who wrote the computer programs giving the results in Figures 4 and 5.

distribution narrows because less uncertainty remains. At the final stage, all uncertainty has been driven out and benefits are known.²

As contrasted with benefits, which are a terminal payout, costs in a SDP are a running payment, additive across stages. The cost of getting from one stage to the next is typically uncertain, but once paid it is a sunk cost.

In a SDP, the reasons for paying running costs to move the project along depend upon the form of the research, development, or exploration that is being undertaken. If the research, say, is necessary in the sense that benefits cannot be received unless the research is completed, then there is the traditional motive, present in purely deterministic situations, of bringing the project closer to completion so that benefits can be realized. However, in an uncertain environment there is the additional motive of paying running costs to obtain more *information* about potential benefits. This may be desirable even if the benefits per se are unaffected by the research activity. For instance, in oil exploration a seismographic test does not affect the net benefit that would be received by constructing an oil rig and extracting the available oil. However, such a test may give a good indication of whether it is profitable to exploit a particular field.

In what follows, we will find it useful to distinguish between two types of SDP. In the first (typically R&D), all stages of the project *must* be completed before benefits can be received; in the second (mineral exploration, marketing a new product), the stages of development are optional in the sense that the SDP could in principle be terminated at any stage and the benefits received—however, premature termination implies the resulting benefits to be received are uncertain. In the sequel, the second type of SDP will be called a Two-Sided SDP (for reasons that will become clear later). When the term SDP is used, henceforth it will be reserved for the first or one-sided type of process.

An essential feature of the environment that we are trying to model is the possibility of continuously reviewing a project's status, based upon the latest information. In deciding whether to continue funding, the aim is to maximize expected benefits minus costs. At each step, an optimal decision rule will indicate whether to continue or terminate, based on current information. With the Two-Sided SDP there is also the additional choice at termination of whether to receive the terminal benefits or back out of the project completely.

If the development process did not have a sequential character—if the possibility of continuous review was eliminated so that only a once-and-for-all decision could be made as to whether benefits should be received—then the optimal decision rule would be simple. For a SDP, the project should be undertaken if and only if expected terminal benefits exceed expected costs (to completion). For a Two-Sided SDP, there is no reason to incur research costs and the project should be immediately undertaken if and only if expected terminal benefits are positive. With sequential decision making, a rational policy maker is given more

²As will be made clear later, it is easy to cope with a situation where there is residual uncertainty that cannot be removed by research.

freedom and it is not surprising that the prospect of continuous review tends to make a project look more attractive. It is important to understand why.

As an extreme example, suppose there is some stage with a low cost, very high variance contribution to the benefits. For example, the use of a seismographic test in oil exploration, mentioned earlier, may be relatively inexpensive to perform but could contribute enormously to narrowing down the uncertainties. Although the expected net benefits from the exploitation of the well might be negative, the optimal policy may still be to commission seismographic tests, and then build a rig if prospects are encouraging or back out if prospects are unpromising. The bias toward tentatively going ahead with a project is more pronounced as the variance of benefits is greater because the realization of a stage removes more uncertainty and allows a better informed decision to be made.

We require a sequential decision rule which indicates whether or not to continue at each stage as a function of the information then available. The optimal stopping rule maximizes expected benefits minus costs, taking account of the fact that at all future stages we will also be following an optimal stopping rule. In the general case, an optimal stopping rule will be very complicated. It will depend upon such things as the underlying probabilistic process, past history, the distribution of costs and benefits, the stage where the project is currently located, etc.

Our primary aim is to derive a simple operational criterion for funding a SDP or a Two-Sided SDP. Such a criterion should depend in a straightforward way on basic parameters of the development project so that it can be easily applied and analyzed. Yet, given this aim, it should not be based upon assumptions or regularity conditions that are unduly restrictive.

We concentrate for the most part on (One-Sided) SDPs where benefits can only be received after all stages have been completed. In Section 2 a model is laid out and the optimal decision rule is presented. Our intention is that this section should be self contained and understandable to the non-specialist reader, even though the underlying mathematics is quite technical. Section 3 is concerned with the properties of the model and the derivation of the optimality criterion. Sections 4 and 5 then deal with a relaxation of the assumptions required for the results of Sections 2 and 3. The Two-Sided SDP is considered in Section 6, and concluding remarks are offered in Section 7.

2. A MODEL OF SEQUENTIAL DEVELOPMENT

We suppose that the SDP must pass through a series of distinct stages whose order is considered to be rigidly prescribed at the outset; further, benefits will only be received if all stages of the project are completed. For example, in an R&D project the first stage might be the construction of a prototype model, followed by intensive development of a particular component, then development of another component, etc. The assumption that a sequential development

project can be well ordered by its stage of development is a mathematical abstraction that allows the entire decision structure to be reduced to a stop-or-go problem at each step. The only choice is between terminating or continuing on to the next step. Thinking of each stage as fulfilling some particular function in the development process, a complete temporal ordering does not seem wholly implausible.

The state variable is the number of steps or stages from the end. At stage s there are s steps remaining to be completed before the project can be implemented and benefits collected. Viewed with the information available at stage s , the final benefit is a random variable, denoted X_s . As s becomes smaller, more information is available and X_s is made progressively less uncertain until, at $s = 0$, X_0 is no longer random.

Development costs for each stage are also uncertain. But by contrast with benefits, development costs are pay-as-you-go or running costs. Once paid, the expenditure of realizing a stage is in effect a sunk cost. If development effort ceases before the project ends, no benefit is received, and viewed *ex post*, the previously sunk development costs have nothing to show. If development continues to completion through all s remaining project stages, the *expected* extra cost is C_s .

To render the economic analysis of a SDP tractible, we need further assumptions. We must postulate the specific form of the probability density function for benefits, as well as the exact relationship between costs, stages of development, and the spread of the distribution of benefits. Lacking such regularity assumptions, we would be faced with a very general process and it would be extremely difficult to obtain a strong characterization of the optimal stopping rule. In Sections 4 and 5 we will discuss what happens in more general cases or with different conditions.

There are three basic regularity assumptions.

ASSUMPTION 1: The random variable X_s , representing terminal benefits as perceived at stage s , is normally distributed for each s .

Of all the specific families of distributions to assume, the normal is perhaps the least objectionable in a context like this. Basically we are assuming that the potential reward, which could only be collected if all research work were completed, can be viewed as a sum of independent random variables over a large number of component development stages. As additional research money is paid to develop another stage, its additive "contribution" to the final reward becomes known. This shifts the mean of potential reward (if the contribution of a stage turns out better or worse than expected) and it narrows the variance because fewer uncertain stages remain to be resolved.

The normality postulate will be re-examined in Section 5. We will show that normality follows from some rather weak assumptions about continuous information gathering.

Suppose at stage s , X_s has a standard deviation σ_s and at stage t , $t < s$, X_t has a standard deviation $\sigma_t < \sigma_s$. Then, *viewed from stage s* : (a) the mean of X_s is known, say μ_s , and

$$(1) \quad X_s \sim N(\mu_s, \sigma_s^2);$$

$$(b) \quad X_t \sim N(m_t, \sigma_t^2) \text{ where } m_t \sim N(\mu_s, \sigma_s^2 - \sigma_t^2).$$

Statement (b) is true because, in order for information to be consistent, $(X_s - \mu_s)$, which is $N(0, \sigma_s^2)$, must be distributed as $(X_t - m_t)$, which is $N(0, \sigma_t^2)$, plus $(m_t - \mu_s)$. It follows from the way independent normal distributions add that $(m_t - \mu_s)$ must be $N(0, \sigma_s^2 - \sigma_t^2)$.

Thus, as the project moves forward, the estimate of terminal benefits contracts in variance, while its mean changes to be consistent with the assumption of normality.³ The decision whether or not to continue at any stage will depend upon (among other things) the expected benefit viewed at that stage, which is changing stochastically as development proceeds.

Note that X_s can be negative with non-zero probability, e.g., the net terminal benefit of doing R&D on an alternative technology could in principle be less than zero if the value or cost reduction of the new technology turned out to be less than that of an existing substitute. The best way to think of X_s is what *would* be obtained *if* the project were forced through all the way to completion from stage s (leaving out development running costs). In principle, a negative net benefit would never be realized or observed because the project could be stopped before its end.

As the SDP passes through stages, the standard error of the distribution of terminal benefits is narrowing down. To obtain neat results, we must postulate that this is occurring in a regular way:

ASSUMPTION 2a: *For any stages s and t*

$$(2) \quad \frac{\sigma_s}{\sigma_t} = \frac{C_s}{C_t}.$$

The standard deviation of estimated terminal benefits at any stage is assumed proportional to the expected cost of carrying the project through to completion from that stage. Assumption 2a means that if the project is developed to a stage where expected remaining costs are half as large, the standard deviation of estimated benefits is also halved. This is a specific way of quantifying the notion that the spread of benefits becomes narrower as development proceeds. It is a very powerful aid in obtaining neat results and does not appear to be grossly objectionable, perhaps because so little is known about how uncertainty is

³As we are restricting ourselves to the family of normal distributions, the variance is an unambiguous measure of the degree of uncertainty.

reduced in most research, development, or exploration projects.⁴ The fact that Assumption 2a permits us to sharply characterize an optimal solution makes it a natural preliminary to any more general analysis. And it may even be a reasonable description of some situations.

In Section 4, a more general process will be considered:

ASSUMPTION 2b:

$$(3) \quad \sigma_s = k(C_s)^\gamma, \quad k, \gamma > 0.$$

Such a formulation embodies hypothesis 2a when $\gamma = 1$ as well as the hypotheses that the standard deviation falls faster than expected costs to completion ($\gamma > 1$) and that the standard deviation falls slower than expected costs to completion ($\gamma < 1$). Actually, in Section 5 it will be shown that a much weaker assumption than even 2b has strong implications:

ASSUMPTION 2c: $\sigma_s = h(C_s)$ for some well-defined, continuously differentiable, monotonically increasing function h satisfying $0 = h(0)$.

The main content of this assumption is that σ_s is systematically related to C_s . We shall have nothing to say about situations where 2c does not hold.

Now that the underlying probabilistic process is well specified, it is easy to state the problem, at least informally (a rigorous statement is temporarily deferred). At any stage the relevant issue is whether to proceed to the next stage or terminate. An optimal stopping rule is defined by backwards recursion. At each stage an optimal stopping rule maximizes expected benefits minus costs based upon information available at that stage and taking account of the fact that an optimal stopping rule will be used for all stages hence, based on the information available at those later stages.

A final regularity assumption is needed to simplify the analysis. So far we have been treating steps as if they were discrete. Now let each step size become smaller and smaller until in the limit s is a continuous variable. This embodies the notion that project review is a continuous undertaking. A stop-go decision can be made at any point because each stage is presumed infinitesimally small relative to total project size.

ASSUMPTION 3: *The stage of development, s , is a continuous state variable.*

The following theorem holds under the assumptions that have been laid down (including Assumption 2a).

THEOREM 1—The Optimal Stopping Rule: *At any stage, let the expected cost to completion of the project be C and the perceived benefit be distributed $N(\mu, \sigma^2)$. It is*

⁴For instance, although Klein [2] has suggested that the standard deviation falls faster than expected costs to completion, Scherer [5], in a comment on that paper, suggested that the reverse proposition is likely in many instances. Assumption 2a is a natural compromise.

optimal to proceed with the SDP if and only if

$$(4) \quad C \leq E[Y | Y \geq \mu]$$

where $Y \sim N(0, \sigma^2)$.

The righthand side of (4) is simple to calculate:

$$(5) \quad E[Y | Y \geq \mu] = \frac{\int_{\mu}^{\infty} yf(y) dy}{\int_{\mu}^{\infty} f(y) dy}$$

$$(6) \quad = \frac{\sigma \exp(-\mu^2/2\sigma^2)}{\sqrt{2\pi} (1 - \Phi(\mu/\sigma))} = \frac{\sigma\Phi'(\mu/\sigma)}{\Phi(\mu/\sigma)}$$

where

$$(7) \quad f(y) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-y^2/2\sigma^2},$$

$$(8) \quad \Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-z^2/2} dz.$$

Perhaps the easiest way of understanding (4) is to ask: given C and σ , what is the cutoff value of μ , call it $\hat{\mu}$, which turns the inequality (4) into an equality? This is defined by

$$(9) \quad C = E[Y | Y \geq \hat{\mu}].$$

At $\mu = \hat{\mu}$, the decision maker would be just indifferent between continuing and stopping the project. When $\mu < \hat{\mu}$, it is best to stop, whereas if $\mu > \hat{\mu}$, the optimal policy is to continue funding.

Using (5), $\hat{\mu}$ must obey

$$(10) \quad \int_{\hat{\mu}}^{\infty} (y - C)f(y) dy = 0,$$

or

$$(11) \quad \int_C^{\infty} (y - C)f(y) dy = \int_{\hat{\mu}}^C (C - y)f(y) dy.$$

Equation (11) allows a simple geometric interpretation of (9), because the $N(0, \sigma^2)$ probability weighted average distance from C to ∞ must equal the weighted average distance from C to $\hat{\mu}$. This is a great help in visualizing the dependence of $\hat{\mu}$ on C and σ , even if it provides insufficient explanation of the optimal rule. Although equations like (10) or (11) are suggestive, the underlying process is complicated and the real justification for the theorem is its proof.

It is very easy to determine the basic properties of the cutoff benefit $\hat{\mu}$: (i) $C > 0$ and $\sigma > 0$ implies that $\hat{\mu} < C$. (ii) $\sigma \rightarrow 0$ implies that $\hat{\mu} \rightarrow C$. (iii) $\sigma \rightarrow \infty$ implies that $\hat{\mu} \rightarrow -\infty$. (iv) $C \rightarrow 0$ implies that $\hat{\mu} \rightarrow -\infty$. (v) $C \rightarrow \infty$ implies that $\hat{\mu} \rightarrow \infty$.

Performing comparative statics on (10) or (11), as σ increases, $\hat{\mu}$ decreases. With greater uncertainty in benefits, the decision maker should be willing to tolerate a lower mean when deciding to continue funding. If σ is larger, it may be optimal to check out a few more stages for superoptimistic outcomes even though μ is low.

As C increases, so does $\hat{\mu}$. The rationale behind this is obvious.

Note that there is nothing in principle to prevent $\hat{\mu} < 0$. When $\hat{\mu} < \mu < 0$, it is optimal to proceed even though the expected value of terminal benefits is less than zero (and this does not take account of running costs!). Less than zero benefits would, of course, never be realized because the process would be discontinued before completion. From (4) and (6), $\hat{\mu} \geq 0$ as

$$\frac{\sigma}{C} \geq \sqrt{\frac{\pi}{2}} \approx 1.25.$$

3. MICRO-ASPECTS OF THE SDP

Although the model we have outlined is well-specified, its form is inappropriate for rigorous analysis. In this section, a mathematical formulation as a Weiner like diffusion process is presented which is equivalent to our original model.

We have already established that, at any given stage, the mean of the distribution of terminal benefits is known with certainty, but is a random variable when viewed from an earlier stage. In other words, the mean of terminal benefits, μ , may be treated as undergoing stochastic shocks. Given C and σ as functions of the number of stages to completion of the process, all relevant information is summarized in the number of stages left to completion, s , and the mean of the (normal) distribution of terminal benefits, μ_s .

Figure 1 shows one possible realization of the SDP. As more stages are completed (we move to the right), the mean μ changes. At $s = 0$, the distribution of benefits has collapsed to μ_0 , the actual reward made possible by the research, development, or exploration.

Figure 1 gives an example where the SDP is pushed through to completion. However, our major interest is in rules which tell us whether, at any stage, the project should be terminated. Such a decision will depend upon available information, i.e., (given the functions $C(s)$ and $\sigma(s)$) the pair $\{\mu_s, s\}$. Two properties of the stop-go decision may be immediately cited:

(i) If $\{\mu_s, s\}$ is such that $\mu_s > C_s$, then the project should be continued. To see this, it may be noted that $\mu_s - C_s > 0$ is the expected return if the sub-optimal rule of always continuing the project is followed, whereas if the project is terminated, the return is zero.

(ii) If $\{\mu_s, s\}$ is such that it is optimal to continue with the project, then it is optimal to continue under $\{\mu, s\}$, $\mu > \mu_s$. This follows because one could




FIGURE 1

'pretend' that expected benefits were μ_s , follow the optimal rule as though starting from μ_s , and then take the *extra* benefit $\mu - \mu_s$ if the project were to be completed. From our normality assumption, the distribution of $\mu_t - \mu_s$, $t < s$, given μ_s , is independent of μ_s , so that increments in the stochastic process are state independent.

From these properties it may be inferred that for all s there exists a $\hat{\mu}(s)$ (possibly $-\infty$) with the property that the project should be continued if $\mu_s > \hat{\mu}(s)$ and discontinued if $\mu_s < \hat{\mu}(s)$.⁵ Our task is to derive the optimal stopping function $\hat{\mu}(s)$.

Given a stopping function, $\hat{\mu}(s)$, it is theoretically possible to calculate the expected net value of the project starting at any $\{\mu_s, s\}$. Let $V(\mu_s, s)$ be the *state valuation function* which gives the value of the project when an *optimal* stopping rule is followed. V can never be strictly negative because a policy of direct termination ensures a net value of zero.

From the definition of $\hat{\mu}(s)$,

$$(12) \quad V(\hat{\mu}(s), s) = 0, \quad \forall s.$$

Further, at $s = 0$ all uncertainty has been removed and V is given by

$$(13) \quad V(\mu, 0) = \max(\mu, 0).$$

Now let us look at the stochastic process governing μ_s . First, as it is being assumed that there is a continuum of stages of development, any reasonable

⁵A continuity argument may be used to show that the decision maker will be indifferent about continuing if $\mu_s = \hat{\mu}(s)$.

normalization concerning s can be adopted. For convenience, *define* the distance s to the end of the project by the fact that expected running costs to the end of the project are equal to s , i.e.,

$$(14) \quad C_s = s.$$

This is an admissible normalization if C_s is declining and positive, e.g., running costs at each stage are positive. As we are assuming that σ_s is proportional to C_s , the normalization imposes the constraint that

$$(15) \quad \sigma_s = s\bar{\sigma}$$

for some constant

$$(16) \quad \bar{\sigma} = \sigma_s / C_s$$

representing the ratio of standard deviation to cost.

Thus the underlying parameter of the system will be given by $\bar{\sigma}$, and the state valuation function V will be

$$(17) \quad V = V(\mu, s; \bar{\sigma}).$$

Although development is being considered as a continuous process, it is useful to visualize the SDP as passing through discrete stages and then taking limits. Assume that each stage is of length l . From stage s to stage $s - l$, the results of development will either be encouraging or discouraging. Encouraging results lead to an upward revision of μ . Assume that the change in μ , $\delta\mu$, is given by

$$(18) \quad \delta\mu_s = Z\bar{\sigma}\sqrt{2sl}$$

where $E[Z] = 0$, $\text{var}[Z] = 1$. The random variable Z gives the distribution of increments; for instance, Z may take on the values 1 and -1 , each with probability $\frac{1}{2}$, and this corresponds to the notion that each stage of the development process is either a success or a failure. The mean of the incremental change is zero because any bias will have been predicted and, therefore, is already embodied in μ_s .

The rationale for (18) is that it yields (15). The variance of terminal benefits at distance s from completion will be equal to the sum of the variances of increments which occur before completion. As at time s there will still be s/l stages to completion, (18) gives

$$(19) \quad \begin{aligned} \text{var}[X_s] &= E \left[\sum_{\tau=1}^{s/l} \delta\mu_{\tau l}^2 \right] \\ &= E \left[2Z^2\bar{\sigma}^2l^2 \sum_{\tau=1}^{s/l} \tau \right] \\ &= \bar{\sigma}^2(s^2 + sl). \end{aligned}$$

Thus, as l tends to zero,

$$(20) \quad \sigma_s^2 = \bar{\sigma}^2 s^2$$

or

$$(21) \quad \sigma_s = \bar{\sigma} s$$

which is exactly (15). Further, as X_s is constructed from a large number of independent increments, the Central Limit Theorem gives

$$(22) \quad X_s \sim N(\mu_s, \sigma_s^2),$$

so that, when l tends to zero (18) is compatible with the assumptions laid down in the last section. In actuality, as l tends to zero (18) becomes a stochastic differential equation (see Cox and Miller [1]):

$$(23) \quad d\mu_s = Z\bar{\sigma}\sqrt{2s} ds .$$

We have seen how the SDP may be viewed as the limit of an additive process. The converse is also true. Under the normality postulate and (15), we may note that (23) must hold. For,

$$\mu_t - \mu_s \sim N(0, s^2\bar{\sigma}^2 - t^2\bar{\sigma}^2)$$

which implies that

$$\begin{aligned} (\mu_t - \mu_s)^2 &= Z^2\bar{\sigma}^2(s^2 - t^2) \\ &= Z^2\bar{\sigma}^2(s + t)(s - t) \end{aligned}$$

where $E[Z] = 0$, $\text{var}[Z] = 1$. Letting $t = s - \delta s$,

$$(\mu_s - \mu_{s-\delta s})^2 = Z^2\bar{\sigma}^2(2s - \delta s)\delta s,$$

so that, if $\delta s \rightarrow 0$,

$$d\mu^2 = Z^2\bar{\sigma}^2 2s ds, \quad \text{i.e., } d\mu = Z\bar{\sigma}\sqrt{2s} ds$$

which is (23).

Our *only* reason for working with continuous rather than discrete stochastic processes is just that it enormously simplifies the solution of the dynamic programming problem.

We are now in a position to investigate the form of the state valuation function V and the stopping curve $\hat{\mu}(s)$. In the mathematics literature, similar problems to ours have been studied under the heading of "Optimal Stopping with Free Boundaries." For a survey of this area, see [5]. Our problem belongs to a class whose general properties, e.g., existence, differentiability, have been examined in [6]. Here, these general properties are taken for granted and the specific form of the relevant functions is investigated.

Consider $\mu_s > \hat{\mu}(s)$. Given the stochastic specification of μ , in a sufficiently small time interval μ_s will not become less than the stopping value $\hat{\mu}$ with probability approaching unity.⁶ Thus, within a small time interval, the stopping rule can be ignored. In such a situation, $V(\mu, s)$ is the expected value of $V(\mu + d\mu, s - ds)$ after the payment of a running cost with an expected value ds :

$$(24) \quad V(\mu, s) = E\left[V(\mu + Z\bar{\sigma}\sqrt{2s ds}, s - ds) \right] - ds.$$

From a Taylor expansion, valid to terms of first order and less in ds ,

$$E\left[V(\mu + Z\bar{\sigma}\sqrt{2s ds}, s - ds) \right] = V(\mu, s) + V_\mu \bar{\sigma}\sqrt{2s ds} E(Z) - V_s ds + V_{\mu\mu} \bar{\sigma}^2 s ds E(Z^2),$$

where subscripts denote partial derivatives. This must hold for all sufficiently small ds , and combined with (24) gives the usual Kolmogorov type diffusion equation:

$$(25) \quad V_{\mu\mu} \bar{\sigma}^2 s - V_s - 1 = 0.$$

Now consider the form of V when $\mu = \hat{\mu}(s)$. If $\hat{\mu}(s)$ is continuous, when the realization of Z is negative, V will be zero at $s - ds$. Thus,

$$(26) \quad V(\hat{\mu}(s), s) = E\left[V(\hat{\mu} + Z\bar{\sigma}\sqrt{2s ds}, s - ds) \mid Z \geq 0 \right] \cdot \text{pr}[Z \geq 0] - ds.$$

From a Taylor expansion, valid to terms of one-half order,

$$(27) \quad E\left[V(\hat{\mu} + Z\bar{\sigma}\sqrt{2s ds}, s - ds) \mid Z \geq 0 \right] \cdot \text{pr}[Z \geq 0] - ds = V(\hat{\mu}, s) + V_\mu \bar{\sigma}\sqrt{2s ds} E[Z \mid Z \geq 0] \cdot \text{pr}[Z \geq 0] + 0(ds),$$

which yields for sufficiently small ds , as $\text{pr}[Z > 0] \neq 0$,

$$(28) \quad V_\mu(\hat{\mu}(s), s) = 0.$$

Equations (12), (13), (25), and (28) should, in principle, completely characterize the form of the function V .

However, we have yet to exploit a homogeneity condition. As V is an expected monetary value, it must be measured in the same units as μ , C_s , and σ_s . V must therefore satisfy

$$(29) \quad V(\mu, s; \bar{\sigma}) = \lambda V(\mu/\lambda, s/\lambda; \bar{\sigma}), \quad \forall \lambda > 0,$$

i.e., V is linearly homogeneous in μ and s . Thus, for fixed $\bar{\sigma}$, V can be written in the functional form

$$(30) \quad V(\mu, s) \equiv \mu g(\mu/s).$$

⁶This also requires the continuity of $\hat{\mu}(s)$ which can be proved rigorously without difficulty.

Now, (12), (13), (25), and (28) become

$$(12') \quad g(\hat{\mu}(s)/s) = 0,$$

$$(13') \quad g(\infty) = 1,$$

$$(25') \quad \bar{\sigma}^2 x g''(x) + (2\bar{\sigma}^2 + x^2) g'(x) - 1 = 0,$$

$$(28') \quad g'(\hat{\mu}(s)/s) = 0.$$

The form of optimal stopping function is easy to analyze. As the state evaluation function is linearly homogeneous, it must be the case that

$$(31) \quad \hat{\mu}(s) = \alpha s, \quad \forall s,$$

for some slope parameter α . (12') and (28') become

$$(32) \quad g(\alpha) = g'(\alpha) = 0.$$

Solving (25') gives

$$(33) \quad g'(x) = \frac{1}{x^2} \left[-k \exp\left(-\frac{x^2}{2\bar{\sigma}^2}\right) + 1 \right]$$

where $k > 0$ is a constant of integration. Using (32) to eliminate k gives

$$(34) \quad g'(x) = \frac{1}{x^2} \left[1 - \exp\left(\frac{\alpha^2 - x^2}{2\bar{\sigma}^2}\right) \right].$$

The differential equation (25') holds when the project is still continuing, i.e., when $x \geq \alpha$. Therefore, it is permissible to write (using (13') and (32))

$$(35) \quad 1 = g(\infty) - g(\alpha) = \int_{\alpha}^{\infty} g'(x) dx$$

(which implicitly defines α). Integrating (34) by parts gives

$$(36) \quad g(\infty) - g(x) = \frac{1}{x} \left[1 - \exp\left(\frac{\alpha^2 - x^2}{2\bar{\sigma}^2}\right) \right] + \frac{1}{\bar{\sigma}^2} \exp(\alpha^2/2\bar{\sigma}^2) \int_x^{\infty} \exp(-v^2/2\bar{\sigma}^2) dv.$$

Using (35),

$$(37) \quad 1 = \frac{1}{\bar{\sigma}^2} \exp(\alpha^2/2\bar{\sigma}^2) \int_{\alpha}^{\infty} \exp(-x^2/2\bar{\sigma}^2) dx.$$

However,

$$(38) \quad \int_{\alpha}^{\infty} x \exp(-x^2/2\bar{\sigma}^2) dx = \bar{\sigma}^2 \exp(-\alpha^2/2\bar{\sigma}^2),$$

so that (37) can be written in the form

$$(39) \quad 1 = \frac{\frac{1}{\sqrt{2\pi\bar{\sigma}^2}} \int_{\alpha}^{\infty} x \exp(-x^2/2\bar{\sigma}^2) dx}{\frac{1}{\sqrt{2\pi\bar{\sigma}^2}} \int_{\alpha}^{\infty} \exp(-x^2/2\bar{\sigma}^2) dx}.$$

Making use of the formula for the normal distribution, this becomes

$$(40) \quad 1 = E[W | W \geq \alpha]$$

where $W \sim N(0, \bar{\sigma}^2)$. Condition (40) defines the slope α of the optimal stopping function. If $\mu_s/s \geq \alpha$, then the project should be continued. Using (14) and (15) to eliminate s , which has been used only as a carrier variable, the project should be continued if

$$(41) \quad C_s < E[Y | Y \geq \mu_s]$$

where $Y \sim N(0, \sigma_s^2)$. This is precisely the stopping rule presented in the last section.

Basic properties of the stopping rule have already been considered and will not be reiterated here. However, one feature of the rule may be noted. Since accumulated information is not useful for determining future running costs, the fact that they are uncertain has no influence upon the decision maker. As soon as running costs are expended, they become a sunk cost.⁷ Thus, there is a fundamental asymmetry in the SDP between running cost uncertainty and terminal benefit uncertainty.

Having obtained the optimal stopping rule, now let us investigate the form of the state valuation function V . Inserting (37), (38), and (39) into (36) yields, after some manipulations,

$$(42) \quad V(\mu, s) = \mu g(\mu/s) = \mu - C_s + \frac{\text{pr}\{Y \geq \mu\}}{\text{pr}\{Y \geq \hat{\mu}(s)\}} [E[Y | Y \geq \mu] - \mu]$$

where $Y \sim N(0, \sigma_s^2)$ and $\hat{\mu}(s)$ is the stopping value of μ at s .

Turning to an analysis of (42), we have already stressed that there exists the possibility of utilizing accumulated information during the course of a SDP. If such information could not be utilized then the expected net benefit of the process would be given by

$$(43) \quad \tilde{V}(\mu, s) = \mu - C_s, \quad \mu \geq C_s.$$

Thus, for $\mu \geq C_s$, the value of being able to utilize accumulated information is

⁷With a different decision criterion, sunk costs could influence optimal behavior. For instance, this would be the case if the objective function was expected utility and the utility function embodied a nonconstant degree of absolute risk aversion.

given by the last term in (42), i.e.,

$$(44) \quad I = \frac{\text{pr}\{Y \geq \mu\}}{\text{pr}\{Y \geq \hat{\mu}(s)\}} [E[Y | Y \geq \mu] - \mu],$$

and (42) can be rewritten as

$$V(\mu, s) = \tilde{V}(\mu, s) + I.$$

I is strictly positive if and only if the variance σ_s^2 is strictly positive. Of course, information will only have value if it is likely to be utilized. It can be shown that $I_{\bar{\sigma}} (= V_{\bar{\sigma}}) > 0$. A situation with higher variance has greater value of information.

It is not surprising that I diminishes with μ (but not at a rate sufficient to diminish V):

$$(45) \quad -1 \leq I_{\mu} = -\frac{\text{pr}\{Y \geq \mu\}}{\text{pr}\{Y \geq \hat{\mu}(s)\}} < 0.$$

Notice also that

$$(46) \quad I_{\mu\mu} > 0$$

so that V is a convex function of μ . The common sense of this result is that an increase in μ has more value the greater the probability of completion of the development process. This probability increases as μ increases, since it is then less likely that the stopping rule will ever be invoked.

To further understand the function V , it is useful to look at V_s . If V_s is positive then a decision maker would like to undertake more “roundabout” development processes (bigger s). This increases the uncertainty concerning terminal benefits, and also makes total running costs greater. In terms of the function g ,

$$(47) \quad V_s = -\left(\frac{\mu}{s}\right)^2 g'(\mu/s)$$

so that the sign of V_s depends upon the sign of g' (the sign of V_s is independent of the particular normalization of s which is adopted). From (34),

$$(48) \quad V_s \geq 0 \quad \text{iff} \quad |\mu/s| \leq |\alpha|.$$

As $\mu/s \geq \alpha$ (V only takes the form in (42) when $\mu \geq \alpha s$), $V_s > 0$ if $\alpha < 0$ and $|\mu/s| < |\alpha|$. These conditions hold if the variance of the process is large enough ($\alpha < 0$ iff $C_s < \sigma_s \sqrt{2/\pi}$) and if the realization of the process is sufficiently close to the stopping line. In this case, the value of being able to exploit accumulated information is so great that there would be a gain from adopting a more roundabout development process even though higher expected costs to completion would be involved.

Figures 2 and 3 give the form of iso- V curves for the two cases $\alpha > 0$ and $\alpha < 0$.




FIGURE 2.—($\alpha > 0$).




FIGURE 3.—($\alpha < 0$).

Two interesting features of these regimes may be noted. First, the slope of the iso- V curves becomes 1 at $s = 0$. This is because when s is close to zero, the probability of non-completion is so small that it may be ignored.

Second, when the process is close to the stopping line, the use of information is very important. If the variance is small relative to costs, so that $\alpha > 0$, the process will be continued so long as “average” outcomes arise, i.e., μ_s stays constant.

However, when $\alpha < 0$, the process will be continued only if “successful” outcomes arise, i.e., μ_s rises. When the variance is large relative to costs, it may be optimal to continue with a process that is unlikely to be completed in the hope that future stages of development will be more successful than is anticipated *a priori*.

4. NON-HOMOGENEOUS SDP'S

In view of the above analysis, it is appropriate to call the process where σ_s is proportional to C_s a homogeneous SDP. We now wish to consider the characteristics of more general processes. A natural class to investigate is obtained by relaxing Assumption 2a of Section 2 to Assumption 2b, i.e., to the class of processes for which

$$(3) \quad \sigma_s = k(C_s)^\gamma$$

for some $k, \gamma > 0$. To take an example, the development of a new aircraft involves wind-tunnel tests on small models. Such tests are likely to significantly reduce σ_s with little reduction in C_s . In this case, γ will exceed unity.

A different pattern emerges if each stage of the micro development process is identical to any other. Assume that expected running costs per stage of development are 1 and that the stochastic differential equation governing μ_s is given by

$$(49) \quad d\mu = Z\bar{\sigma}\sqrt{ds}$$

where $E[Z] = 0, \text{var}[Z] = 1$.

The importance of (49) is that $d\mu$ is independent of s . Thus, (49) will be recognized as a Wiener process with zero drift. The macro-characteristics of the process are easily analyzed:

$$(50) \quad C_s = \int_0^s ds = s$$

and

$$(51) \quad \sigma_s^2 = E\left[\int_0^s Z^2\bar{\sigma}^2 ds\right] = \bar{\sigma}^2s.$$

Thus,

$$(52) \quad \sigma_s = \bar{\sigma}\sqrt{s} = \bar{\sigma}\sqrt{C_s},$$

so that the process belongs to the class given by (3) with $\gamma = \frac{1}{2}$.

In fact, $\gamma = \frac{1}{2}$ is a natural lower bound in many circumstances. If development takes the form of passing through a large number of additively independent stages, then $\gamma < \frac{1}{2}$ implies that stages with highest variance per unit expected running cost occur at the end of the project. However, should the *order* of stages be subject to choice, stages with the highest variance per unit running cost will be

completed first. For it is optimal to try to know whether the project should be completed after the smallest possible outlay of running costs. This is achieved by first undertaking those stages of development which give the most information per unit expected running cost.

The determination of an explicit formula for $\hat{\mu}(s)$ seems to be a difficult task when $\gamma \neq 1$. For this reason, only the general shape of the stopping curve will be investigated. Not surprisingly, the homogeneous case where $\gamma = 1$ turns out to be a dividing line between two qualitatively different regimes.

As in the last section, s may be normalized so that

$$(53) \quad C_s = s.$$

In this case

$$(54) \quad \sigma_s = ks^\gamma.$$

We start with two simple properties of the optimal stopping curve $\hat{\mu}(s)$. First, $\hat{\mu}(0) = 0$. Second, $\hat{\mu}(s) \leq s$ as, when $\mu > s = C_s$, a policy of always completing the project yields positive expected net gain.

Now assume that the development process has reached a point $\{\mu, s\}$. It is clear that $\mu > \hat{\mu}(s)$ if a positive expected net gain would result from a policy of pushing the project through all stages and then checking whether the (now certain) terminal benefits from the project should be realized. The expected net gain from such a policy is given by

$$(55) \quad \hat{V} = E[Z | Z \geq 0] - s$$

where $Z \sim N(\mu, k^2s^{2\gamma})$. Assume that $\mu = \nu s$. Simple manipulation gives

$$(56) \quad \hat{V} = \left[ks^{(\gamma-1)} E \left[W | W \geq \frac{-\beta s^{(1-\gamma)}}{k} \right] + \beta - 1 \right] s$$

where $W \sim N(0, 1)$.

Consider $\gamma < 1$ first. As $s \rightarrow 0$, we have

$$(57) \quad \hat{V} \rightarrow kE[X | X \geq 0]s^\gamma > 0$$

so that, whatever the value of ν , it is optimal to continue when s is close enough to zero, i.e., for $\gamma < 1$,

$$(58) \quad \left. \frac{d\hat{\mu}}{ds} \right|_{s=0} = \infty.$$

Next, consider $\gamma > 1$. As $s \rightarrow \infty$, (57) holds. Thus, whatever the value of ν , it is optimal to continue when s is large enough, i.e., for $\gamma > 1$,

$$(59) \quad \frac{\hat{\mu}(s)}{s} \rightarrow -\infty \quad \text{as } s \rightarrow \infty.$$

The investigation of limits in the other direction requires a different approach.

Although it can be established rigorously, let it be assumed that $\hat{\mu}(s)/s$ goes to a limit as s tends to zero or infinity. When $\gamma < 1$, $\mu^* = \lim_{s \rightarrow \infty} \hat{\mu}(s)/s$ must be investigated. It is already known that $\mu^* \leq 1$. It will be shown that if $\{\mu_s, s\}$ is such that $\mu_s/s = \nu$ where $\mu^* < \nu < 1$ then $V(\nu s, s)$ must be negative for sufficiently large s . This contradicts optimality and so it must be the case that $\mu^* = 1$.

Only the general argument will be sketched. Assume that $\mu/s = \nu$ where $\mu^* < \nu < 1$. If s is sufficiently large, the decision rule says that the project should be continued. We can ask the question: should the project be continued at least $\tau s^{(1+\gamma)/2}$ stages, for some fixed τ ? It is not difficult to show that the distribution of the change in μ , $\delta\mu$, over this time period is given by

$$(60) \quad \delta\mu \sim N(0, 2\gamma k^2 \tau s^{(5\gamma-1)/2}).$$

Thus the change in μ/s , $\delta(\mu/s)$, is given by

$$(61) \quad \delta(\mu/s) = Z(2\gamma k^2 \tau)^{1/2} s^{(5/4)(\gamma-1)} - \nu \tau s^{(1/2)(\gamma-1)}$$

where $Z \sim N(0, 1)$. With $\gamma < 1$, the distribution collapses to a point as $s \rightarrow \infty$. Thus, with probability approaching unity, μ/s will not reach μ^* and so the project will be continued for at least $\tau s^{(1+\gamma)/2}$ with an expected net running cost of $\tau s^{(1+\gamma)/2}$.

Next consider the expected net gain, given that the process starts at $\{\nu s + \delta\mu, s - \tau s^{(1+\gamma)/2}\}$. It is clear that the expected net gain cannot exceed the expected net gain that would result if the actual path of μ_s is known *ab initio*. If the future is known, only those paths which lead to a benefit which exceeds running costs will be followed. Expected net gain will be given by

$$(62) \quad V = [E[W | W \geq \tilde{s}] - \tilde{s}] \text{pr}\{W \geq \tilde{s}\}$$

where $W \sim N(\nu s + \delta\mu, k^2 \tilde{s}^{2\gamma})$ and $\tilde{s} = s - \tau s^{(1+\gamma)/2}$. Manipulation of this expression yields, when $s \rightarrow \infty$,

$$(63) \quad V \leq \frac{ks^\gamma}{\sqrt{2\pi}}.$$

We can now state that $V(\nu s, s)$ cannot exceed V after the payment of a running cost with expected value $\tau s^{(1+\gamma)/2}$. Using (63),


$$(64) \quad V(\nu s, s) \leq \frac{ks^\gamma}{\sqrt{2\pi}} - \tau s^{(1+\gamma)/2}$$

when $s \rightarrow \infty$. Thus, as $s \rightarrow \infty$,

$$(65) \quad V(\nu s, s) \leq s^\gamma \left[\frac{k}{\sqrt{2\pi}} - \tau s^{(1-\gamma)/2} \right] < 0.$$

This is the contradiction that we wished to obtain: it must be the case that

$$(66) \quad \hat{\mu}(s)/s \rightarrow 1 \quad \text{as } s \rightarrow \infty.$$

FIGURE 4.—($\gamma = \frac{1}{2}$).

A similar analysis applies when $\gamma > 1$ and $s \rightarrow 0$, i.e., for $\gamma > 1$,

$$(67) \quad \left. \frac{d\hat{\mu}}{ds} \right|_{s=0} = 1.$$

These limiting properties give a good indication of the general shape of the optimal stopping curve $\hat{\mu}(s)$. Figures 4 and 5 give the results of numerically solving the two cases $\gamma = \frac{1}{2}$ and $\gamma = 2$ by computer. In these figures, the axes are scaled so that when $\sigma_s = C_s$, values of s , σ_s , and C_s are all set at unity. These curves exhibit the qualitative features just obtained and are intuitively plausible. For instance, if $\gamma > 1$ and s is large, the completion of more stages of development significantly reduces the variance of terminal benefits at little cost. It is not surprising that in this case a very weak decision rule is applied at first.

The results for $\gamma > 1$ are in accord with those derived for the homogeneous SDP. If the process $\gamma > 1$ starts at s , the average value of σ_t/C_t , $t \leq s$ begins at zero when $s = 0$ and then increases monotonically to infinity at $s = \infty$. The analysis of the homogeneous SDP suggests that for the case $\gamma > 1$, $\hat{\mu}(s)/s$ should begin at 1 when $s = 0$ and decrease to $-\infty$ at $s = \infty$. Our investigation of the limits of the non-homogeneous process $\gamma > 1$ show that this is actually true. An analogous interpretation would show that results for the case $\gamma < 1$ are also in accord with the homogeneous SDP.

Finally, we have so far been maintaining that the distribution of terminal




FIGURE 5.—($\gamma = 2$).

benefits shrinks to a point as s tends to zero. However, the analysis is in no way affected if the random variable of terminal benefits is given by

$$(68) \quad \chi_s = X_s + R$$

where X_s and R are independent and the mean of R is zero. The uncertainty associated with R is uncertainty that cannot be reduced by research, development, or exploration. As should be clear, the distribution of R does not alter the decision rules that have been discussed, the effect of R being eliminated when expected values are taken.

5. THE NORMALITY ASSUMPTION AND THE THEORY OF INFORMATION GATHERING PROCESSES

So far, the assumption that the distribution of benefits is normally distributed has been rigidly maintained. The purpose of this section is to consider under what conditions normality follows, which requires a slightly more general study of information gathering processes.

Abandoning the axiom of normality, we shall retain the assumption that the SDP is a continuous process. Rather than presuming that the relationship between σ_s and C_s is of a particular form, Assumption 2c will be invoked, i.e.,

$\sigma_s = h(C_s)$ where h is some well-defined continuously differentiable, monotonically increasing function with $0 = h(0)$. Normalizing s so that $C_s = s$, Assumption 2c implies that σ_s is a well-defined function of s , which means that the expected cost of uncertainty reduction at stage t is known at stage s , $s > t$. This is a powerful restriction because, along with an assumption of the form that small amounts of information can only change beliefs by a small amount, it will imply the normality postulate.

Let r_s be the information that has accrued by the time that the SDP reaches stage s . In general, a decision maker's distribution of terminal benefits at stage s will depend upon r_s . Let this distribution function be $f_s(x | r_s)$. Between stages s and t , $s > t$, further information, ${}_s r_t$, will be received and, at stage s , there will be a probability measure on ${}_s r_t$, call it H , which may depend upon accumulated information r_s . The following relationship will hold:

$$(69) \quad f_s(x | r_s) = \int_{{}_s r_t} f_t(x | r_t) dH(r_s).$$

Define

$$(70) \quad \mu_s(r_s) = \int_x x f_s(x | r_s) dx$$

to be the mean of the distribution of terminal benefits as viewed at stage s . Since at stage s , ${}_s r_t$ is unknown, $\mu_t(r_t)$ will be a random variable. Consider

$$(71) \quad \Delta\mu(r_s) = \mu_t(r_t) - \mu_s(r_s)$$

which is also a random variable. We have

$$(72) \quad E[\Delta\mu(r_s)] = \int_{{}_s r_t} \mu_t(r_t) dH(r_s) - \mu_s(r_s).$$

Making use of (69) and (70),

$$(73) \quad E[\Delta\mu(r_s)] = 0.$$

Given this, the variance is

$$(74) \quad \text{var}[\Delta\mu(r_s)] = \int_{{}_s r_t} (\mu_s(r_s) - \mu_t(r_t))^2 dH(r_s).$$

Now, we know that

$$(75) \quad \int_x (x - \mu_s(r_s))^2 f_s(x | r_s) dx = \sigma_s^2$$

for all s and r_s . Using (75), manipulation of (74) yields

$$(76) \quad \text{var}[\Delta\mu(r_s)] = \sigma_s^2 - \sigma_t^2$$

which, by assumption (2c), is independent of r_s .

The remarkable feature of (73) and (76) is that the first two moments of the

distribution of the difference between expected terminal benefits at two future stages in the SDP are independent of information that is received during the course of undertaking the SDP previous to the two stages.

The usefulness of $\Delta\mu(r_s)$ is that $\mu_0(r_0)$ is, as $h(0) = 0$, the terminal benefit that may be received. By (71), the distribution of $\Delta\mu(r_s)$ when $t = 0$ gives the distribution $f_s(x | r_s)$, displaced so as to make the mean zero. Thus, to understand the distribution of f_s , it is sufficient to understand the distribution of $\mu_0(r_0)$ when viewed at s .

If the process governing μ is Markovian, then normality of f follows directly. Without the Markovian assumption normality still follows from a weak condition. Specifically, it may be reasonable to assume that *small amounts of information can only change beliefs by a small amount*, i.e., μ_s is a *continuous* process. If this is the case then (73) and (76) imply that μ is a square-integrable martingale. This together with (76) implies, by a theorem of Doob, that $\mu_0(r_0)$ when viewed at s is a normal distribution or, more importantly, the distribution functions f_s are normal distributions.⁸

The analysis of this paper is therefore relevant when Assumption 2c is invoked and a continuity assumption is imposed. Although it is easy to comprehend these assumptions, it is useful to describe situations where they are not satisfied:

(i) The essence of Assumption (2c) is that at stage s , it is possible to know how much uncertainty will be left (σ_t^2) at stage t , $t < s$. Thus the uncertainty reduction brought about by moving the SDP forward is known with certainty. If Assumption (2c) is abandoned then an extra degree of uncertainty is added to the model. Under the continuity assumption, this uncertainty about uncertainty reduction must always be present when distribution functions are non-normal.

(ii) When continuity does not hold, then continuing research can lead to a sudden breakthrough.⁹ For instance, over a short-interval l , μ may make jumps which are independent of l but with probability proportional to l . The first two moments of the distribution will still conform to (73) and (76) but continuity will not hold. In this case, the distribution functions f will be given by a Poisson-like distribution which will approach normality only as s becomes large. In general, of course, SDP's can embody both continuous and discrete changes in beliefs but an investigation of such processes is well beyond the scope of this paper.

6. THE TWO-SIDED SDP

We have, so far, analyzed a process where all stages of development must be completed before benefits are received. In that model, research and/or development plays the double role of moving the process closer to the stage where terminal benefits are received, and providing information concerning the value of terminal benefits.

⁸For the appropriate definitions, see Lipster and Shirayev [3]. For Doob's theorem, see their Theorem 5.12. We are indebted to M. Rothschild for pointing out that the result proved in an earlier version of the paper was but a special case of the Doob theorem.

⁹An apt terminology might be the Eureka-SDP.

The purpose of this section is to briefly consider the Two-Sided SDP, where only the second role exists. That is, the completion of development stages provides information concerning the value of terminal benefits but, at any stage, it is possible to move directly to the completion of the process. Thus a dual decision is faced at each stage: the development process can be terminated or continued; given termination, either the project which provides the terminal benefits may be undertaken or it can be discarded.

An example of a Two-Sided SDP is a firm considering marketing a new product when consumer reaction is uncertain. At any moment the firm could decide to market the product, or not to market it at all; or, the firm might, at a cost, perform further product testing, conduct market surveys, and the like.

Many irreversible decisions in natural resources or environmental management can be viewed as Two-Sided SDP's. Examples could be drawn from several other areas. Choosing one of two alternative uncertain projects can frequently be modeled as a Two-Sided SDP.

In this section we retain the assumptions of the homogeneous SDP (continuity, normality, proportionality of σ_s and C_s). Equation (29) still applies so that the state valuation function, V , may be written in the form

$$(30) \quad V(\mu, s) \equiv \mu g(\mu/s).$$

There will now exist two stopping curves. It is clear that there must exist a $\hat{\mu}(s)$ such that if $\mu \leq \hat{\mu}(s)$ then the process should be terminated and the final project should not be undertaken. Similarly, there must exist a $\tilde{\mu}(s)$ ($\tilde{\mu}(s) \geq \hat{\mu}(s) \forall s$) such that if $\mu \geq \tilde{\mu}(s)$ the process should be terminated and the final project should be undertaken. From the homogeneity assumption, these stopping curves will be linear, i.e.,

$$(77) \quad \begin{aligned} \hat{\mu}(s) &= \beta s, \\ \tilde{\mu}(s) &= \zeta s, \end{aligned}$$

where $\zeta \geq \beta$. Such a configuration is shown in Figure 6. The reason for adopting the Two-Sided SDP terminology is now clear.

Equations (12') and (28') apply as earlier and the differential equation (25') continues to hold when μ_s lies between $\hat{\mu}(s)$ and $\tilde{\mu}(s)$. We must investigate the form of V at the stopping line $\tilde{\mu}(s)$. If the development process realizes a point $\{\mu, s\}$ where $\mu \geq \tilde{\mu}(s)$ then the terminal benefits will be taken directly. By definition, the expected reward associated with such an enterprise is given by μ , i.e.,

$$(78) \quad V(\mu, s) = \mu, \quad \mu \geq \tilde{\mu}(s).$$

Applying this result in (24) gives

$$(79) \quad V_\mu(\tilde{\mu}(s), s) = 1.$$




FIGURE 6

As $V_\mu = g + (\mu/s)g'$, (78) and (79) become

$$(80) \quad g(\zeta) = 1$$

and

$$(81) \quad g'(\zeta) = 0.$$

$g'(x)$ is given by (33) so that (28') and (81) yield

$$(82) \quad 1 - k \exp\left(-\frac{\beta^2}{2\bar{\sigma}^2}\right) = k \exp\left(-\frac{\zeta^2}{2\bar{\sigma}^2}\right).$$

Thus, $|\beta| = |\zeta|$. However, as $g(\beta) \neq g(\zeta) \rightarrow \beta \neq \zeta$, $\zeta \geq \beta$ implies that

$$(83) \quad \zeta \geq 0$$

and

$$(84) \quad \beta = -\zeta.^{10}$$

Thus the stopping lines must be symmetrical about the horizontal axis in

¹⁰Although $\beta = \zeta = 0$ is possible, notice that g would fail to be differentiable at zero.

Figure 6. To characterize the slope parameters, we may note that

$$(85) \quad 1 = g(\zeta) - g(\beta) = \int_{-\zeta}^{\zeta} g'(x) dx,$$

so that, pursuing a similar analysis to that in Section 3, we obtain

$$(86) \quad 1 = E[Z | Z \geq \zeta] \left[\frac{\text{pr}\{Z \geq \zeta\}}{1 - 2\text{pr}\{Z \geq \zeta\}} \right]$$

where $Z \sim N(0, \bar{\sigma}^2)$.

Performing comparative statics, ζ goes up when $\bar{\sigma}$ increases or when s decreases. As $\bar{\sigma}$ tends to zero, or as s tends to infinity, ζ tends to zero. As $\bar{\sigma}$ tends to infinity, or s tends to zero, ζ tends to infinity.

As might have been suspected, the symmetrical feature noted in (84) is a general result for Two-Sided SDP's. Figure 7 gives the basic shape of the stopping curves for the cases studied earlier where $\sigma_s = k(C_s)^\gamma$.

Note that μ can be interpreted as the difference between the means of the distributions of benefits associated with two projects, one of which must be chosen. Interpreting μ as $\mu_2 - \mu_1$, "termination and discontinuation" means selecting project 1, "termination and undertaking" means selecting project 2. In this context, running costs will be costs expended to both evaluate and discrimi-




FIGURE 7

nate between the two projects, as well as the implicit cost of delaying both projects for a period. With the interpretation of a choice between two projects, it is easy to understand why stopping curves must be symmetric.

As an example, it may have been decided that an electricity generating station is to be built and the decision is whether it should be coal or nuclear powered. The Two-Sided SDP may be used to derive rules which tell a planner that either more research into the net benefits associated with the two options should be undertaken, or one of the plant designs should be adopted. From our analysis it may be noted that, independent of research costs, in an optimal policy more information should be sought if both options seem to offer the same expected benefits, given any uncertainty.

It may be remarked that the Two-Sided SDP is indicative of the way in which the results and methodology of earlier sections can be extended to cover other related situations.

7. CONCLUDING REMARKS

This paper has examined optimal sequential decision rules for research, development, and exploration projects. Instead of summarizing the analysis, we will, in concluding, simply note the fact that a natural taxonomy of processes has been uncovered.

First, we have seen that it is important to distinguish between those processes for which the stages of development must be completed before benefits can be received and those where development stages are optional. Most real-world processes contain stages belonging to both categories although sometimes the dichotomy is clear cut, e.g., market research contains only optional stages.

Second, the analysis of Section 5 suggests that it is useful to distinguish between processes where small amounts of information only change beliefs by small amounts and processes where there are sudden breakthroughs or setbacks. The form of the distribution functions which capture uncertainty of terminal benefits depends crucially upon this feature of the development process.

Finally, it has been shown that there is a qualitative difference between those processes for which it is the case that remaining running costs fall faster than the standard deviation of terminal benefits and those processes for which it is the case that remaining running costs fall slower than the standard deviation of terminal benefits.

Our hope is that a categorization of real-world processes along these lines, perhaps accompanied by some crude calculations based on the formulae of this paper, will provide a helpful start toward understanding the rational application of funding criteria.

St. Catherine's College, Oxford
and
Massachusetts Institute of Technology

REFERENCES

- [1] COX, D. R., AND H. D. MILLER: *The Theory of Stochastic Processes*. London: Chapman and Hall, 1965.
- [2] KLEIN, B. H.: "The Decision Making Problem in Development," in *The Rate and Direction of Inventive Activity: Economic and Social Factors*, National Bureau of Economic Research. Princeton: Princeton University Press, 1962.
- [3] LIPSTER, R. S., AND A. N. SHIRYAYEV: *Statistics of Random Processes*. Heidelberg: Springer-Verlag, 1977.
- [4] SCHERER, F.: "Comment on Klein," in *The Rate and Direction of Inventive Activity: Economic and Social Factors*, National Bureau of Economic Research. Princeton: Princeton University Press, 1962.
- [5] VAN MOERBEKE, P. L. J.: "Optimal Stopping and Free Boundary Problems," *Rocky Mountain Mathematics Journal*, 4(1974), 539-578.
- [6] ———: "On Optimal Stopping and Free Boundary Problems," *Archives For Rational Mechanics and Analysis*, 60(1976), 101-148.