NBER WORKING PAPER SERIES

OPTIMAL DYNAMIC R&D PROGRAMS

Gene M. Grossman

Carl Shapiro

Working Paper No. 1658

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 1985

We are greatly indebted to Avinash Dixit for his invaluable input into Sections II and III, and to Robert Willig for his suggestions that allowed us to extend our earlier results in Section IV. We thank Michael Katz, David Newbery, and seminar participants at the University of Pennsylvania for their comments, and the National Science Foundation under Grant #SES-8408622 for financial support. Grossman also acknowledges the support of the Alfred P. Sloan Foundation. The research reported here is part of the NBER's research program in Productivity. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

NBER Working Paper #1658
July 1985

Optimal Dynamic R&D Programs

ABSTRACT

We study the optimal pattern of outlays for a single firm pursuing an R&D program over time. In the deterministic case, (a) the amount of progress required to complete the project is known, and (b) the relationship between outlays and progress is known. In this case, it is optimal to increase effort over time as the project nears completion.

Relaxing (a), we find in general a simple, positive relationship between the optimal expenditure rate at any point in time and the (expected) value at that time of the research program. We also show that, for a given level of expected difficulty, a riskier project is always preferred to a safe project. Relaxing (b), we find again that research outlays increase as further progress is made.

Gene M. Grossman
Woodrow Wilson School
Princeton University
Princeton, NJ 08544
(609) 452-4823

Carl Shapiro
Woodrow Wilson School
Princeton University
Princeton, NJ 08544
(609) 452-4847

I. Introduction

Many research projects, as well as some types of investment programs for the installation of physical capital, can be described as follows: measurable progress is achieved over a period of time, but the investment yields no returns until the entire project is completed. Examples of this include laboratory development of a new product or process, the construction of a new building, and the writing of a scholarly journal article. When confronted with investment opportunities of this sort, individuals and firms must decide how many resources to devote to the project at each point in time. Implicitly, this also determines the (expected) duration of the project.

In this paper we characterize the optimal time path of R&D outlays when a "prize" is earned only after some discrete amount of progress is achieved. We study both deterministic R&D ventures, where the amount of progress necessary for success is known at the outset, and uncertain or "risky" projects, where the difficulty of the endeavor is initially unknown. In the latter case advancement will imply not only the completion of more stages of the research, but also an updating of beliefs about the distance still to be covered. Our characterization will include a comparison of the dynamics of safe and risky research endeavors.

In the analysis that follows we treat the determination of a dynamic R&D investment profile as an optimal control problem facing a single firm. This can be thought of as an uncontested pursuit of a patent for a new product, or a project to improve the technology for producing output in a competitive industry. (The dynamics of oligopolistic interaction in a multi-phase patent race are the topic of our current research efforts.) We present the solution to the deterministic program in Section II. In Section III, we analyze a model of risky R&D in which the relationship between expenditures and progress is known, but the amount of progress necessary to complete the project is unknown. We discuss the effects of this type of uncertainty on both the value of the program and the optimal expenditure pattern in Section IV. Then, in Section V, we study a rather different formulation of riskiness: the amount of progress required for completion is known, but the relationship between effort and progress is stochastic. A final section summarizes our findings.

The remainder of this introduction places our work in context in relation to the existing R&D literature. In the game-theoretic literature, research and development is most often modelled either as a static allocation problem (e.g., Dasgupta and Stiglitz (1980)), or as a dynamic problem only in the limited sense that there is a flow probability of success at any point in time that depends on the current level of effort (e.g., Lee and Wilde (1980)). In the latter case the decision environment is static until an innovation is made, so there is no reason for a firm to alter its behavior over time. Thus, in the most common theoretical formulations of the determinants of R&D outlays, the question of the dynamic program of spending does not arise.

More recent contributions to this literature have emphasized the investment-like qualities of expenditures on R&D. Fudenberg, et.al. captured the notion of "progress" in their model of a two-firm patent race. However, throughout most of their paper they did not allow firms to vary the intensity of research effort, which of course precluded study of the profile of R&D spending.[1] Harris and Vickers (1985) also modelled progress in their recent

---

[1]In a final section of their paper, Fudenberg, et.al. (1983) do allow firms to choose among two levels of effort, but they are concerned with a different set of issues than we address here, namely how the presence of information lags regarding a rival's actions affects competitive racing strategy.

paper on patent races. They specified a (deterministic) research project facing each of two firms that is quite similar to the one that we posit for our single firm in Section II. Furthermore, they found that in a perfect equilibrium only one firm actually engages in research, and it "almost always" acts as it would if it faced no rivalry. However, they did not investigate the investment behavior of this single firm, as they concentrated instead on the determinants of the identity of the winning firm. Finally, Judd (1985) has studied a patent-race model that bears some similarity to the stochastic version of our control problem in Section III, but he was able to provide results only under the restrictive assumption that the prize for success is negligibly different from zero.[2]

Our work is most closely related to two earlier papers in the decision-theoretic R&D literature. Lucas (1971) studied the same deterministic R&D problem that we analyze in Section II. We extend his characterization of the optimal research program and provide an analysis of its comparative dynamics. Kamien and Schwartz (1971) have studied one case of the risky R&D investment program that we analyze in Section III. We use different techniques which allow us both to characterize more clearly and fully the solution to the control problem and to give an economic interpretation of the first-order conditions. Our approach then permits us to study the effects of riskiness on R&D outlays in Section IV.

---

[2]The model in Reinganum (1981) shares some of the features of our formulation, but as we discuss below, the dynamics there are generated entirely by the artifact of an assumed terminal date before which time all research projects must be completed.

II. Optimal R&D Programs for Projects of Known Difficulty

A firm seeks a prize of size W.[3] To obtain this prize it must "travel" a distance L. We denote by x_t the distance that has been covered by date t with $x_0 = 0$. Progress is achieved via the expenditure of resources. Let the rate of advance be given by $\dot{x}_t = f(c_t)$, where c_t is the R&D outlay at time t. We assume that there exists a $\bar{c} \geq 0$ such that $f(c) = 0$ for all $c \leq \bar{c}$, and $f'(c) > 0$ and $f''(c) < 0$ for all $c \geq \bar{c}$. In other words, we posit decreasing returns to effort at any point in time given that some progress is being made, but we allow for the possibility of a fixed start-up cost $\bar{c}$ at every moment.

The firm discounts future receipts and expenditures at rate r. Its problem is to choose expenditures at every point in time up to some terminal date, as well as the date of termination, T, to maximize the presented discounted value of the stream of net profits, subject to the constraint that the total progress attained by the termination date be sufficient to complete the project. We write this control problem as follows:

$$\max_{\{c_t\},T} \; We^{-rT} - \int_0^T c_t e^{-rt} dt$$

$$\text{subject to} \quad \int_0^T f(c_t)dt \geq L$$

---

[3]The amount W can represent the present discounted value <u>at the time of completion of the project</u> of a stream of profits from then into the future. If the firm is risk neutral (which we will assume to be the case below when we introduce uncertainty about the difficulty of the project), then it can also be the expected value of a prize of unknown size.

Let $\pi(W, L)$ be the solution, i.e., the present discounted value of maximal profit when a prize W is at a distance L.

The program is solved in two stages. First we consider the sub-problem for a given terminal date T, the maximized value of which we denote $U(T; W,L)$. The Lagrangian for this sub-problem is

$$H = We^{-rT} - \int_0^T c_t e^{-rt}dt + \lambda[\int_0^T f(c_t)dt - L].$$

The first-order condition for c_t is[4]

$$\frac{\partial H}{\partial c_t} = -e^{-rt} + \lambda f'(c_t) = 0$$

i.e.,

$$f'(c_t) = e^{-rt}/\lambda \tag{1}$$

or

$$c_t = f'^{-1}(e^{-rt}/\lambda) \tag{1'}$$

Substituting (1') into the first-order condition for λ gives

$$\int_0^T f(f'^{-1}(e^{-rt}/\lambda))dt = L \tag{2}$$

---

[4]The second-order condition is $f''(c_t) < 0$, which is satisfied at any (interior) solution. If it is optimal to undertake the research program, then it cannot ever be optimal to choose $c_t = 0$. Doing so would simply delay the entire program, reducing its value.

Let the value of λ that satisfies (2) be written as $\hat{\lambda}(T, L)$ and note for later reference that $\hat{\lambda}_L > 0$ and $\hat{\lambda}_T < 0$, where subscripts denote partial derivatives.

Equation (1) describes a set of paths of R&D expenditures (indexed by λ) that increase over time in such a way that the marginal product of effort falls exponentially at rate r. Lucas (1971) noted this Hotelling-like property for a general f(c) and explicitly solved the problem for the special case $f(c) = c^{\frac{1}{2}}$. Only one of the paths satisfying equation (1) reaches L at time T; higher values of λ imply greater effort at each date. The path that does reach L at time T is optimal for the sub-problem; it has initial effort given by $f'(c_0) = 1/\hat{\lambda}(T, L)$.

Now the solution to the full problem is found by substituting for the arbitrary T above the date of termination that maximizes U(T; W,L). The first-order condition is

$$U_T(T;\ W,L) = -rWe^{-rT} - c_Te^{-rT} + \hat{\lambda}(T,L)f(c_T) = 0$$

or

$$rW = \hat{\lambda}(T,L)f(c_T)e^{rT} - c_T \tag{3}$$

where we have omitted terms involving $\partial c_t/\partial T$ and $\partial\hat{\lambda}/\partial T$ by application of the envelope theorem.[5] The optimal path, $\{c_t\}$, completion time, T(W,L), and corresponding multiplier, $\lambda(W,L)$, simultaneously satisfy conditions (1), (2) and (3).

---

[5]The second-order condition, $U_{TT}(T;\ W,L) < 0$, is satisfied at any interior point where $U_T = 0$. If the project is <u>not</u> worth undertaking, then we will have $U_T > 0$ for all T.

The qualitative properties of the optimally-designed research program can be understood intuitively as follows. If the discount rate is strictly positive, it cannot be optimal to apply effort to a project evenly throughout. Relative to this allocation, the discounted development costs could be decreased for <u>any</u> given duration of the project by shifting expenditures from the early stages to those later on. Only if the discount rate is zero will it be optimal to spend at a constant rate, that rate being the one that maximizes the rate of progress per dollar spent, $f(c)/c$.[6]

Given that the intertemporal pattern of expenditures maintains a constant <u>discounted</u> marginal product of research outlays, the initial intensity of effort is chosen so that the marginal benefit of completing the project a moment sooner is equal to the marginal cost of doing so. The marginal benefit is $rW + c_T$, the instantaneous return on holding the prize plus the expenses that would no longer need to be borne at T if the project were to be completed before then. The marginal cost is $f(c_T)/f'(c_T)$, the extra outlay that would be required to travel a distance $f(c_T)$ farther at the moment before time T so as to ensure completion of the project. Note that $rW + c_T = f(c_T)/f'(c_T)$ is implied by equation (3), after substitution of the first-order condition for c_T from equation (1).

It is straightforward to see how the optimal path varies with the parameters of the problem. In Figure 1, we have plotted R&D intensity as a function of the stage of the project. Relative to some base path, an increase in the

[6] It is clear from (1) that when $r = 0$, expenditures are time invariant. Substituting $f'(c_T) = 1/\hat{\lambda}$ into (3) gives $f'(c_T) = f(c_T)/c_T$, which is the first order condition for maximizing $f(c_T)/c_T$. With no fixed costs of R&D, $c_t \to 0$ as $r \to 0$. When fixed costs are present ($\bar{c} > 0$) there will be some $c^* > \bar{c}$ that maximizes $f(c_T)/c_T$. Effort is constant at this level when there is no discounting, and always greater than this level otherwise.

prize implies that the project should be completed sooner and that initial effort should be greater. This is because

$$\partial T/\partial W = -U_{TW}/U_{TT} = re^{-rT}/U_{TT} < 0$$

and

$$\frac{\partial c_0}{\partial W} = -\frac{1}{\hat{\lambda}^2 f''(c_0)} \left(\frac{\partial \hat{\lambda}}{\partial T}\right)\left(\frac{\partial T}{\partial W}\right) > 0 \ .$$

Substituting $\lambda = e^{-rT}/f'(c_T)$ from equation (1) into equation (3), the terminal R&D intensity, c_T, is given implicitly by

$$rW = \frac{f(c_T)}{f'(c_T)} - c_T \equiv \psi(c_T). \qquad (4)$$

An increase in the prize, W, causes terminal intensity to increase, since $\psi'(c) > 0$. Indeed, R&D outlays are higher at every stage of the project and at every point in time when the prize is bigger. Note finally that profits are a convex function of the size of the prize, i.e., $\pi_{WW} = -re^{-rT(W,L)}(\partial T/\partial W) > 0$. We shall use this fact in our analysis of riskiness below.

An increase in the difficulty of the research project, L, raises the time to completion. This can be seen by using $U_L = -\hat{\lambda}$ and $\hat{\lambda}_T < 0$ to give

$$\frac{\partial T}{\partial L} = \frac{-U_{TL}}{U_{TT}} = \frac{1}{U_{TT}} \ \frac{\partial \hat{\lambda}}{\partial T} > 0.$$

Terminal effort is given by (4) and therefore is unaffected by changes in L.

Most important for our purposes is the fact that the value of the research program is convex in its difficulty, L. In other words, $\pi_{LL}(W,L) = -\lambda_L(W,L) > 0$, i.e., $\lambda_L < 0$. To establish this point, note that $\pi(W,L_1 + L_2) = \pi(\pi(W,L_2), L_1)$, i.e., a prize of W at distance $L_1 + L_2$ is equivalent to a prize of $\pi(W,L_2)$ at distance L_1. Differentiating this identity with respect to L_1 gives $\pi_L(W,L_1+L_2) = \pi_L(\pi(W,L_2),L_1)$. Differentiating next with respect to L_2 gives $\pi_{LL}(W,L_1 + L_2) = \pi_{LW}(\pi(W,L_2),L_1)\pi_L(W,L_2)$. Since $\pi_L < 0$, all that remains to be shown is that $\pi_{LW} < 0$. But $\pi_{LW} = - \lambda_W(W,L)$ and $\lambda(W,L) = \hat{\lambda}(T(W,L),L)$, so $\pi_{LW} = - \hat{\lambda}_T(\partial T/\partial W)$. Finally, $\hat{\lambda}_T < 0$ and $\partial T/\partial W < 0$, which implies $\pi_{LW} < 0$ as claimed.

We have established that, as the difficulty of the project increases, the marginal cost of having to travel yet farther, $\lambda(W,L)$, falls. In view of equation (1), this implies that when the project is more difficult research is undertaken less vigorously at every stage (and at every moment) up to the last.

As a final comparative dynamics point, the effects of an increase in the discount rate are found by differentiating equations (2) and (4) and again noting the relationship between c_0 and λ from (1). An increase in r reduces initial outlays but increases terminal effort. The net effect on the duration of the research project is ambiguous.

III. Optimal R & D Programs for Projects of Unknown Difficulty

Suppose that the progress required to complete the project initially is unknown, but that effort, c, still leads to progress according to the deterministic function $\dot{x} = f(c)$. The amount of progress needed to attain the prize is assumed to be a random variable L with probability density function p(L) on the support $(\underline{L}, \bar{L})$, and corresponding cumulative distribution function P(L). We define the hazard rate, $\phi(x)$, so that $\phi(x)dx$ is the probability that success will be achieved between x and x+dx, conditional on the project already having progressed (unsuccessfully) to a distance x. Then $\phi(x) \equiv p(x)/(1-P(x))$.

We assume that the firm is risk neutral. Its problem, which is to maximize expected discounted profits, can be written as

$$\max_{\{c_t\}} \int_{x=\underline{L}}^{x=\bar{L}} [We^{-rt(x)} - \int_0^{t(x)} c_\tau e^{-r\tau} d\tau] p(x)dx$$

where t(x) is defined implicitly by $\int_0^{t(x)} f(c_\tau)d\tau = x$.[7] Kamien and Schwartz (1971) have analyzed this problem using Pontryagin methods under the assumption that all R&D activity must end by some exogenously given time.

This control problem is solved most easily and transparently using the techniques of dynamic programming. Define V(x) to be the maximum value function, i.e., the maximal expected discounted profits earned by a firm that has progressed to x and behaves optimally thereafter. The firm's expenditure rate when at x, c(x), will be chosen such that, during the time interval from t to $t+\Delta t$,

$$V(x) = \max_c [-c\Delta t + We^{-r\Delta t}\phi(x)f(c)\Delta t + (1-\phi(x)f(c)\Delta t)V(x+\Delta x)e^{-r\Delta t}]$$

[7] An alternative formulation of an R&D project of unknown difficulty is one where "success" requires a "breakthrough," the probability of which is an increasing function of current effort and cumulative past progress, say $\rho(c,x)$. If $\rho(c,x)$ is assumed to be separable and of the form $f(c)\phi(x)$, where f(c) is also the measure of advancement that enters into the definition of x (i.e., $x_t = \int_0^t f(c_\tau)d\tau$), then this formulation is equivalent to ours.

where we have used $\Delta x = f(c)\Delta t$. The first term in the brackets is the direct cost of the research program. The second term is the discounted value of the prize times the probability of completion of the project during the interval under consideration. This latter probability is the product of the hazard rate, $\phi(x)$, and the amount of ground covered, Δx. The final term is the probability that success is not attained during the interval multiplied by the discounted value of the program at the end of the interval in the event that this is the case.

For small time intervals, we have $V(x+\Delta x) \cong V(x) + V'(x)\Delta x$ and $e^{-r\Delta t} \cong 1-r\Delta t$. Also, terms in $(\Delta t)^2$ and $(\Delta t)^3$ vanish in the limit as $\Delta t \to 0$. Substituting, and taking this limit, we have

$$V(x) = \max_c \, [-c + \phi(x)f(c)W - V(x)\phi(x)f(c) - rV(x) + f(c)V'(x)]\Delta t + V(x)$$

or

$$rV(x) = \max_c \, [-c + f(c)\{\phi(x)(W-V(x)) + V'(x)\}] \tag{5}$$

From this we compute the first-order condition for $c(x)$,

$$f'(c(x)) = \frac{1}{\phi(x)(W-V(x)) + V'(x)} \tag{6}$$

We assume for the moment that equation (6) has a solution, and that it is optimal to continue the research program when at x. Finally, we insert the optimized R&D outlay from equation (6) into the Bellman equation (5), to obtain

$$rV(x) = \frac{f(c(x))}{f'(c(x))} - c(x) \equiv \psi(c(x)) \tag{7}$$

Note the parallel between equations (7) and (4).

Equation (7) gives the optimal research expenditure at x (assuming that it is non-zero) as a function of the value of the program; that is, $c(x) = \psi^{-1}(rV(x))$. Since $\psi'(c) > 0$, (7) implies that effort is greater when the <u>current</u> position is highly valued.[8] We now can substitute for c(x) in (6), to arrive at the following differential equation describing V(x):

$$V'(x) = g(V(x)) + \phi(x)(V(x)-W) \tag{8}$$

where

$$g(V) \equiv \frac{1}{f'(\psi^{-1}(rV))}$$

If the value function found by solving (8) subject to the boundary condition $V(\bar{L}) = W$ remains everywhere non-negative, then this gives the solution to the dynamic program. Alternatively, if no such path exists, then it will be optimal either to never begin doing research or to abort the research program if success is not achieved by some critical stage.

In the following two subsections we distinguish two alternative situations for purposes of describing the qualitative properties of the optimal program. These are (A) when the hazard rate $\phi(x)$ is everywhere non-decreasing, and (B) when the hazard rate declines for some range of values of x. Subsequently, we discuss the optimal program under two sets of conditions that are excluded by the assumptions maintained in the analysis thus far. These extensions are to

---

[8]Note that if r = 0 the optimal program again involves a constant level of effort, c*, so as to maximize f(c)/c. Henceforth, we shall assume that the discount rate is positive.

(C) the case of a constant hazard rate over an unbounded support for x and (D) the case of a discrete distribution of potential prize locations.

A. Hazard Rate Everywhere Non-Decreasing

Suppose that $\phi'(x) \geq 0$ for all $x \in [\underline{L},\bar{L}]$. This situation arises if whenever a success not realized, researchers become more optimistic that a breakthrough is imminent. This case is the one that Kamien and Schwartz have studied using different methods. Our results in this subsection parallel theirs.

In Figure 2 we plot the set of points in (x,V) space such that $V'(x) = 0$. These points satisfy $\phi(x) = g(V)/(W-V)$. Note that $g'(V) > 0$, which implies that the curve is flat whenever $\phi'(x) = 0$ (including all $x < \underline{L}$) and slopes upward elsewhere. The curve terminates at the point $(\bar{L},W)$, since as $x \to \bar{L}$, $\phi(x) \to \infty$ and we must have $V \to W$ for $V'(x) = 0$ to hold.

We use the $V'(x) = 0$ schedule to aid us in drawing the set of paths that satisfy (8). Above this curve V is rising, while below it V is falling. Two situations are possible. First, there may not exist any path that has $V(x) \geq 0$ for all x and $V(x) \to W$ as $x \to \bar{L}$.[9] In this case, it is not worthwhile to begin the research project at all. Alternatively, there may exist a unique saddlepath that has V(x) everywhere non-negative and that satisfies the boundary condition (as depicted in the figure). Then the saddlepath gives the maximum value function for the dynamic program.

The qualitative properties of the optimal R&D program can now be described. First, note that if it is optimal to begin the research project, then

[9] A necessary condition for this to occur is that for some $x \in [\underline{L},\bar{L}]$ the solution to $V'(x) = 0$ from equation (8) would require $V < 0$. At this x it must be the case that $\phi(x) < 1/Wf'(c^*)$, where we recall that c^* maximizes $f(c)/c$.

it will be optimal to follow it through to fruition in all contingencies. This is because at each moment the prospects for success are at least as bright as they were the moment before. Next, notice that along the saddlepath $V(x)$ is monotonically increasing. But with $r > 0$, research effort varies directly with the value of the program, by equation (7). Thus, if the hazard rate is non-decreasing and the discount rate is positive, the optimal dynamic program involves rising R&D outlays over time.

B. Hazard Rate Declining for a Range of x.

The hazard rate may decline for a range of x before rising again as $x \to \bar{L}$. In the range where $\phi'(x) < 0$, the $V'(x) = 0$ schedule is downward sloping. Three types of outcomes can characterize the optimal program. We illustrate these in the three panels of Figure 3.

In Figure 3a, the optimal program is qualitatively the same as when the hazard rate is everywhere non-decreasing. Although the prospects for success decline for a range of x, the period is not so prolonged nor is the news so negative that it causes the value of the program to fall. Research effort is increasing throughout. This is not the case in Figure 3b. Here, effort is increasing over the stages of research prior to $\underline{L}$. Somewhere before the hazard rate begins to decline it is optimal to begin to reduce R & D expenditures in the event of no success. The process is reversed at a later stage, again prior to the time that the hazard rate begins to increase. If there are several non-contiguous ranges of declining hazard rates, then research intensity may (but need not) change directions several times, with each reversal leading the changes in the sign of $\phi'(x)$.

Finally, Figure 3c depicts a case where it is worthwhile to begin the research program and proceed as far as $\hat{x}$. If success is not achieved by then, it is optimal to abandon the effort. The relevant path for $V(x)$ in this case is the one that has $V(x) = 0$ at exactly the point where the $V'(x) = 0$ schedule meets the horizontal axis. Thus, the shutdown point is given by $f'(c^*)W = 1/\phi(\hat{x})$. This can be seen as follows. Any path above this one ultimately violates the constraint that $V(x) \leq W$. (Note that we are assuming as a precondition for this case that there does not exist a path with $V(x) \geq 0$ for all x and $V(x) \to W$ as $x \to \bar{L}$.) A path that starts below this one can be shown to lead to a contradiction. Consider for example the path that has $V(\tilde{x}) = 0$ at $\tilde{x} < \hat{x}$. With $V'(\tilde{x}) < 0$, direct reasoning would suggest that it is not worthwhile to carry on the program beyond $\tilde{x}$, since the value of the program would turn negative. However, $\tilde{x} < \hat{x}$ and $\phi'(\tilde{x}) < 0$ together imply $\phi(\tilde{x})f'(c^*)W > 1$. Using the definition of c^*, we have $\phi(\tilde{x})f(c^*)W > c^*$. But then it cannot be optimal to cease operations at $\tilde{x}$, because the flow (expected) benefit of proceeding for another instant at intensity c^* exceeds the flow cost. Thus, the path under consideration is not internally consistent.

If a research program has alternating stages where the hazard rate is increasing and decreasing, then it may be optimal for research intensity to change directions several times before a shutdown point ultimately is reached. In any event, research will never cease during a range of increasing hazard rates, and any shutdown point x must satisfy $f'(c^*)W = 1/\phi(x)$.

C. Constant Hazard Rate

For any probability density function with a finite support, the hazard rate eventually must increase as $\bar{L}$ is approached. However, if the support on

$p(L)$ is unbounded, it is possible for the hazard rate to be everywhere constant. This occurs if the density function is exponential, i.e., it has the form $p(L) = \phi e^{-\phi L}$. Reinganum (1981) has investigated a model of R&D activity where this assumption is adopted explicitly. Furthermore, the familiar specification of Lee and Wilde (1980), which posits a flow probability of a breakthrough that depends only on current expenditures, has an alternative interpretation in which effort contributes to "progress", but with a hazard rate on successful completion of the project that is constant (see footnote 7 above).

It is straightforward to extend the analysis of this section to the case of a constant hazard rate. Equation (8), describing the movement of the maximum value function as progress is achieved, continues to apply, but $V(\bar{L}) = W$ must be replaced by an economically-meaningful boundary condition. It is clear that $V(x)$ is bounded above by W and below by zero. Any path that has $V(x) > W$ is logically inconsistent, while $V(x) = 0$ with $V'(x) < 0$ implies that shutdown is optimal. In Figure 4 we show the $V'(x) = 0$ schedule for the case of a constant hazard rate; it is given by $g(V) = \phi(W-V)$, a horizontal line in (x,V) space. Any path that begins above or below this line eventually must violate one of the inequality constraints. Thus, the only consistent path for $V(x)$ has $V'(x) = 0$ everywhere, i.e., the value of the program is constant. It follows that R&D effort is constant as well.[10] This makes intuitive sense, of course, since with a constant hazard rate the problem is completely stationary: failure at any stage does not alter the decision calculus for the future.

---

[10] The optimal level of effort is given implicitly by the solution to $(rW + c)f'(c) = f(c) + r/\phi$. It is increasing in both ϕ and r.

Reinganum (1981) specifies a constant hazard rate and finds as the outcome for a collusive industry that R&D outlays are increasing over time. Kamien and Schwartz (1971) also allow for this possibility, and draw a similar conclusion. There is no contradiction with our result, however, because they assume, unlike here, that all research projects must be completed before an exogenous time T, or else the prize vanishes. Reinganum interprets T as the time that the innovation becomes "obsolete". However, this interpretation is not entirely consistent with another of her assumptions, namely that the (current) value of the prize is independent of the time of discovery.

D. Discrete Distributions

The techniques we have employed thus far in this section are most suitable for control problems where the probability of success of a research project is a continuous function of the extent of progress. The case where p(L) is a discrete distribution is also of some interest, because many research endeavors involve a sequence of discrete experiments, any one of which may yield a discovery. Then the prize can be located only at those points corresponding to the completion of one of these experiments.

We analyze the optimal R&D program for the case of a two-point distribution.[11] Let p and 1-p be the probabilities that the prize is located at distances L_1 and L_2, respectively. When point L_1 is reached, all uncertainty will be resolved. If it turns out that the news at L_1 is unfavorable, the firm will then follow the optimal path associated with a deterministic problem for a prize of W at distance L_2-L_1. Doing so requires

[11] It will become clear that the method extends easily to the n-point case.

that the same level of effort be applied as a function of the stage of the project as would have been optimal between L_1 and L_2 if the firm had known that the prize was located at L_2 from the outset. We illustrate this in Figure 5.

At the instant before L_1, the firm realizes that the (expected) value of the program there is $pW + (1-p)\pi(W,L_2-L_1)$, the probability weighted average of the profits it will earn for each of the two possible outcomes of the initial experiment. Since the firm is risk neutral, it will behave during the stages prior to reaching L_1 exactly as it would under a deterministic regime with a prize of $pW + (1-p)\pi(W,L_2-L_1)$ at a distance L_1. This involves less intensive research at every point than if the prize had been at L_1 with certainty, but greater outlays than if the prize were at L_2 for sure. Research effort is increasing over time within each of the two phases of the project (i.e., during each experiment), but evidently a "disappointment effect" causes a discrete drop in intensity if the initial experiment fails to achieve results.

Finally, we have drawn in Figure 5 for purposes of comparison the path of R&D effort for a deterministic program with a prize W at a distance $pL_1 + (1-p)L_2$. As we shall prove in the next section when we compare the research programs under safe and risky conditions, R&D outlays are everywhere greater during the initial phases of the uncertain regime than they would be if the same prize had been at the mean distance with certainity.

IV. The Effects of Uncertainty on R&D Outlays

In this section we show that a risk neutral firm always prefers a risky research project to a safe project requiring on average the same amount of progress for success. Since research effort is an increasing function of the current value of any program, this will imply immediately that R&D outlays during the stages before the lower support of the distribution of an uncertain regime is reached will exceed those expended at these same stages under a mean-equivalent certain regime. Furthermore, we will show that the effort paths for these alternative programs must have a single crossing, as depicted in Figure 6.

The key to validating these claims is the fact, demonstrated in Section II above, that, for a deterministic research program, profits are a convex function of the distance to be covered: $\pi_{LL}(W,L) > 0$. The convexity of $\pi(W,L)$ in L tells us immediately that if the firm were choosing between a deterministic project of difficulty L and a risky project with expected difficulty L <u>and</u> if the actual difficulty of the risky project were to be revealed before the program had begun (but after the risky project had been selected), then the firm would select the risky option.

Since the actual difficulty of the risky project is learned only over time and after resources have been expended, we <u>cannot</u> conclude from $\pi_{LL} > 0$ that a mean-preserving spread in the p(x) distribution raises V(0). Indeed, we shall present an example below where a mean-preserving spread <u>decreases</u> the value of the program. Nonetheless, any risky regime with expected difficulty μ is always preferred to a safe project with known difficulty μ.

To prove this fact, we begin by comparing the program for a two-point distribution with that for a mean-equivalent safe regime. When the first point is reached where the the prize might be located under the risky option (L_1), the value of this program is $V_R(L_1) = pW + (1-p)\pi(W,L_2-L_1)$. At this same point the value of the program for the safe project is $V_S(L_1) = \pi(W,(1-p)(L_2-L_1))$. It follows immediately from $W = \pi(W,0)$ and $\pi_{LL}(W,L) > 0$ that $V_R(L_1) > V_S(L_1)$. Since optimal paths are followed in each case at all stages prior to L_1 and no learning takes place under either regime during this period, it must be the case that the risky project also is preferred at the outset, i.e., $V_R(0) > V_S(0)$.

The proof extends to distributions with more than two points by backward induction. At the second-to-last point where the prize might be located,[12] say L_{n-1}, the risky project under consideration is preferred to another hypothetical project which is otherwise the same except that, rather than having positive probabilities of discovery of p_{n-1} at L_{n-1} and p_n at L_n, it has a probability $(p_{n-1} + p_n)$ of success at $(p_{n-1}L_{n-1} + p_nL_n)/(p_{n-1} + p_n)$. This latter project in turn has greater value at L_{n-2} than a project with no probability weight there (or at L_{n-1} or L_n), but instead a probability $(p_{n-2} + p_{n-1} + p_n)$ of success at $(p_{n-2}L_{n-2} + p_{n-1}L_{n-1} + p_nL_n)/(p_{n-2} + p_{n-1} + p_n)$. And so, by a chain of inequalities, we have that the initial risky project has higher value at L_1 than a project with the prize located for sure at $\Sigma_i\, p_iL_i$. It must, therefore, have a higher value at the start of the program as well.

We can also show that the paths of optimal effort (as functions of distance travelled) for the safe and risky projects, $c_R(x)$ and $c_S(x)$, cross exactly once. First note, using equation (7), that at any intersection point $V_R(x) = V_S(x)$. The equations of motion for the respective value functions are

$$V_R'(x) = g(V_R(x)) + \phi(x)(V_R(x) - W)$$

and

[12]A continuous probability density function is attained as a limiting case of the n-point discrete distribution as $n \to \infty$.

$$V_S'(x) = g(V_S(x))$$

so that at any point where the values are equal, $V_S'(x) > V_R'(x)$. Thus, once the paths intersect for the first time we must have $V_S(x) \geq V_R(x)$ ever after. Such a first crossing must occur, because $V_R(\underline{L}) > V_S(\underline{L})$, but $V_S(\mu) = W > V_R(\mu)$.

What is the effect on profits of further increases in risk? We demonstrate by means of an example that a mean preserving spread in p(L) may actually lower the value of the optimal R&D program, despite $\pi_{LL} > 0$. That is, although some risk is always better than none, <u>more</u> risk is not necessarily beneficial to the firm. The reason is that some information is obtained earlier under the less risky regime (as well as some later) and in certain circumstances this may allow the firm to readjust its program so as to conserve substantially on (discounted) expenditures.

<u>Example</u>: Let $r = 0$ and c^* minimize $c/f(c)$, and call $\sigma = c^*/f(c^*)$. In all circumstances it is optimal to set c equal to c^* (see equation (7)). Therefore $\pi(W,L) = W - \sigma L$, so long as $W > \sigma L$.

Now consider two risky regimes. Under the riskier regime, A, the prize is at $x = 0$ with probability ¼, at $x = L$ with probability ¼, and at $x = M$ with probability ½, where M is very large. Under the relatively safe regime, B, the prize is at $x = L/2$ with probability ½; otherwise it is at $x = M$. Regime A is a mean preserving spread of regime B.

Consider first regime A. Assuming that $W < \sigma(M-L)$, it will be optimal to quit if the prize is not at $x = L$. Therefore, the value of project A is $W/4 + 3(W/3 - \sigma L)/4$, where we assume that $W > 3\sigma L$. The first term represents the ¼ chance that the prize is at $x = 0$; the second term represents the 1/3 chance that it is at L, given that it is not at 0, which happens with probability 3/4. So, the value of project A is $W/2 - 3\sigma L/4$.

Project B has value $\pi(W/2, L/2) = W/2 - \sigma L/2$, since the expected prize at distance L/2 is W/2. Note that if the prize is not found at L/2 the firm will abandon its efforts there, and again not proceed on to M.

Clearly Project B, the less risky project, is more valuable. The reason is that, if the prize is at the distant point M the firm learns this unfortunate news earlier, and thus can save resources by aborting the project at L/2 rather than at L.

While more risk is not always beneficial for the reason just outlined, it will be so in any situation where initially $c'(x) \geq 0$ for all x in the range of possible prize locations affected by the mean preserving spread.[13] This proposition is proven as follows.[14]

Let $\tilde{c}(t)$ be the program (now expressed as a function of time) that is optimal for p(L) before it is subjected to a mean preserving spread. Define

$$G(L) \equiv We^{-rT} - \int_0^T \tilde{c}(t)e^{-rt}dt$$

where T is given by $\int_0^T f(\tilde{c}(t))dt = L$. Thus G(L) is the profit that is realized if $\tilde{c}(t)$ is followed and it happens that the prize is located at L. We proceed to show that $\tilde{c}'(T(L)) \geq 0$ is sufficient for G(L) to be convex at L.

[13] Any mean preserving spread of p(L) that raises $\bar{L}$ (the upper support of the distribution) is not covered by the conditions of this proposition, since effort falls to zero at $\bar{L}$ in the initial program.

[14] We are grateful to Robert Willig for suggesting this result and its proof to us.

Differentiating G(L) twice gives

$$e^{rT}G''(L) = -[rW + \tilde{c}(T)]\frac{d^2T}{dL^2} + [r^2W + r\tilde{c}(T) - \tilde{c}'(T)](\frac{dT}{dL})^2$$

where $dT/dL = 1/f(\tilde{c}(T))$ and $d^2T/dL^2 = f'\tilde{c}'/f^3$. After substituting, we have

$$f^2e^{rT}G''(L) = \frac{f'\tilde{c}'}{f}[rW + \tilde{c} - \frac{f}{f'}] + r^2W + r\tilde{c}.$$

But $\tilde{c}(t)$ is optimal for the initial program, so $f/f' - \tilde{c}(T) = rV(x(T))$ which is less than rW, and the bracketed expression is positive. Thus, $\tilde{c}'(T) \geq 0$ implies that $G''(L) > 0$. This in turn implies that an increase in risk in a range where effort is initially non-decreasing would raise profits, even if the plan of expenditures were to remain unaltered. It follows _a fortiori_ that profits must increase if a new _optimal_ path is chosen.

Evidently, the adverse effect of increased risk, namely the delay in the revelation of some information, cannot be so severe in cases where R&D expenses are increasing in the relevant range so as to offset the benefit stemming from the fact that the prize _may_ come sooner (which is valuable, since $\pi_{LL} > 0$). The early arrival of unfavorable information under a less risky regime has relatively greatest value to the firm when it plans to reduce its efforts as a consequence.

V. R&D Programs with Stochastic Progress

So far we have assumed that the relationship between current effort and progress is deterministic: $\dot{x} = f(c)$. In this section we study a model of stochastic progress that is a natural generalization of the formulations used in the patent race literature (e.g., Lee and Wilde (1980) and Reinganum (1981)).

Suppose that n distinct tasks $i = 1,2,\ldots,n$, must be completed in sequence to receive the prize, W. The firm's progress is measured by the number of tasks, i, that it has successfully completed. Progress is achieved only by completing the current task; there is no learning while work is underway on a given task. When the firm is engaged in task i, it can achieve a flow probability q_i of completing that task if it spends resources at a rate $c(q_i)$; $c(0) = 0$, $c' > 0$ and $c'' > 0$. Call V_i the value of having completed i-1 tasks, i.e., of being in a position to undertake research on the i^{th} task, $i = 1, \ldots, n$.

We work backwards beginning with $i = n$. Using dynamic programming techniques, if the firm chooses to achieve a completion probability of q, then we have

$$rV_n = -c(q) + q(W-V_n). \tag{9}$$

This equation states that the rate of return r on the asset V_n equals the follow of "dividends," $-c(q)$, plus the expected capital gain, $q(W-V_n)$. Solving for V_n and optimizing the choice of q gives

$$V_n = \frac{q_n W - c(q_n)}{r + q_n} = \max_q \frac{qW - c(q)}{r + q}.$$

So long as $W > \min_q c(q)/q$, it is optimal to select $q_n > 0$. A similiar line of reasoning implies that

$$V_i = \max_q \frac{qV_{i+1} - c(q)}{r + q}, \qquad i = 1, 2, \ldots, n.$$

and

$$q_i = \operatorname{argmax}_q \frac{qV_{i+1} - c(q)}{r + q}, \qquad i = 1, 2, \ldots, n, \tag{10}$$

where we have adopted the labelling convention $V_{n+1} = W$. Naturally,

$$V_i = \frac{q_i V_{i+1} - c(q_i)}{r + q_i} < V_{i+1} .$$

But $\operatorname{argmax}_q \{(qV - c(q))/(r+q)\}$ is increasing in V, so we immediately have $q_1 < q_2 < \ldots < q_n$, i.e., it is optimal to work harder as more progress is made.

From equations (9) and (10), we can express q_n implicitly by $c'(q_n) = W - V_n$; more generally, we have $c'(q_i) = V_{i+1} - V_i$, $i = 1, 2, \ldots, n$. Since $q_{i+1} > q_i$ and $c'' > 0$, this immediately tells us that $V_{i+1} - V_i > V_i - V_{i-1}$ for $i = 2, \ldots, n$. The value of completing the i^{th} task increases with i, even though all the tasks have been assumed to be equally difficult to complete (the same $c(q)$ function applies to all the tasks).

The optimal program with stochastic progress shares two properties with the deterministic program of Section II and the program with unknown difficulty and an increasing hazard rate of Section III. First, it is optimal to increase effort as more stages of the research are completed. Second, each increment of progress is more valuable than was the previous one, i.e., the value of the program is convex in the amount of progress that has been made.

VI. Conclusions

We have identified several reasons why a firm that is optimally pursuing a research and development program will wish to vary the intensity of its efforts over time. All of these reasons rely on the presence of discounting.

First, absent any uncertainty, the firm will find it optimal to work more and more intensively as it comes closer to completing the project. This deterministic result may actually have greater application to conventional investment projects than to research projects *per se*.

Second, and more generally, the firm's optimal intensity of effort should always be an increasing function of the current value of the project itself. If there is a relatively great likelihood that the project soon will be completed, then it is optimal to work relatively hard. If, however, bad news arrives, i.e., it is learned that the project is much more difficult than was previously believed, then it may be optimal to scale back one's efforts, as the value of the project has been diminished.

If the uncertainty facing the researcher regards the amount of progress that will be required for completion, then a risky project where an expected amount of progress L is needed for success will always be preferred to a safe project requiring L for sure. This is so despite the fact that information about the risky project's difficulty will only be revealed over time and after resources have been expended. Further increases in risk only are guaranteed to raise the value of an R&D project if research expenditures are everywhere non-decreasing under the initial program in the range of prize locations affected by the mean preserving spread. In general, more risk can be harmful because it postpones readjustment in situations calling for a sharp reduction in effort.

If the researcher's uncertainty arises from a stochastic relationship between current effort and progress, we again find that it is optimal to increase the level of effort as further advancement is achieved. In this case bad news (no actual progress as a result of yesterday's efforts) simply leaves the researcher in the same position as he was yesterday, so that optimal effort should be unchanged; good news, on the other hand, leads to an increase in the optimal rate of expenditures.

Our analysis in this paper has been limited to the case of a single firm - either a monopolist or a perfect competitor. It is not immediatly evident whether the presence of direct rivalry between firms would accentuate the general principle that has emerged here, namely the tendency of optimal research efforts to increase as progress is made and as time goes by. To the extent that a firm is induced to redouble its efforts when it falls behind, we could have increasing efforts over time in a race. This pattern would also occur if competition intensified when a firm that was formerly behind suddenly caught up. It may be the case, however, that one firm's progress early in the race causes its rival to slow down (or give up entirely); this would be a reason for intense effort early on. We are currently studying R&D rivaly with progress in order to sort out these various effects.

References

Dasgupta, P. and J. Stiglitz, (1980) "Industrial Structure and the Nature of Innovative Activity," Economic Journal, 90, p. 266-293.

Fudenberg, D., Gilbert, R., Stiglitz, J., and Tirole, J. (1983) "Preemption, Leapfrogging and Competition in Patent Races," European Economic Review, 22, pp. 3-31.

Harris, C. and Vickers, J. (1985), "Perfect Equilibrium in a Model of a Race," Review of Economic Studies, 52, pp. 193-209.

Judd, K. (1985), "Closed-Loop Equilibrium in a Multi-Stage Innovation Race," Discussion Paper No. 647, Management Economics and Discussion Science, Kellogg Graduate School of Management, Northwestern University.

Kamien, M. and N. Schwartz (1971), "Expenditure Patterns for Risky R and D Projects," Journal of Applied Probability, 8, pp. 60-73.

Lee, T. and Wilde, L. (1980), "Market Structure and Innovation: A Reformulation," Quarterly Journal of Economics, 94, pp. 429-436.

Lucas, R. (1971), "Optimal Management of a Research and Development Project," Management Science, 17, p. 679-697.

Reinganum, J. (1981), "Dynamic Games of Innovation," Journal of Economic Theory, 25, pp. 21-41.

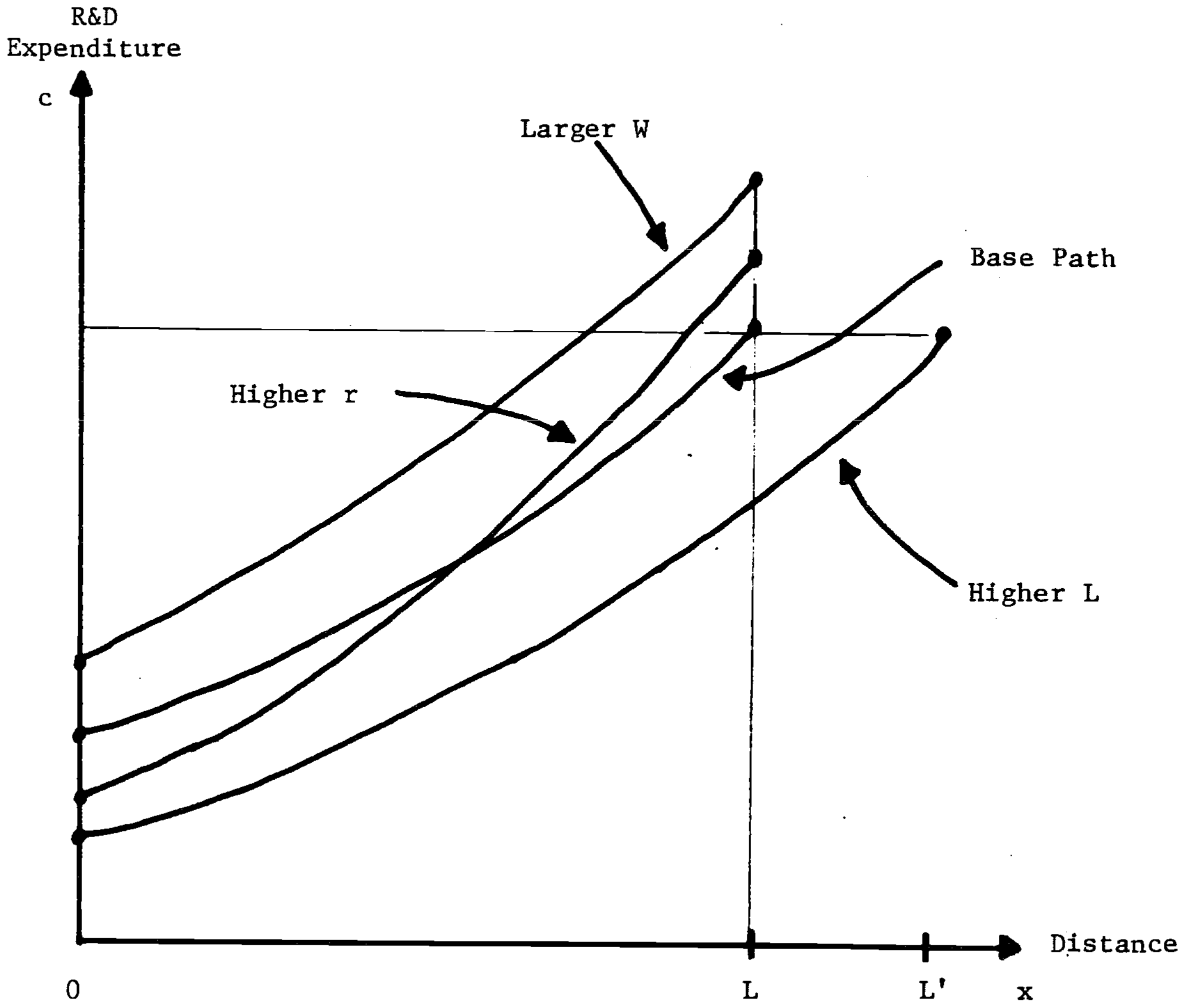

Optimal Expenditure Paths

Figure 1

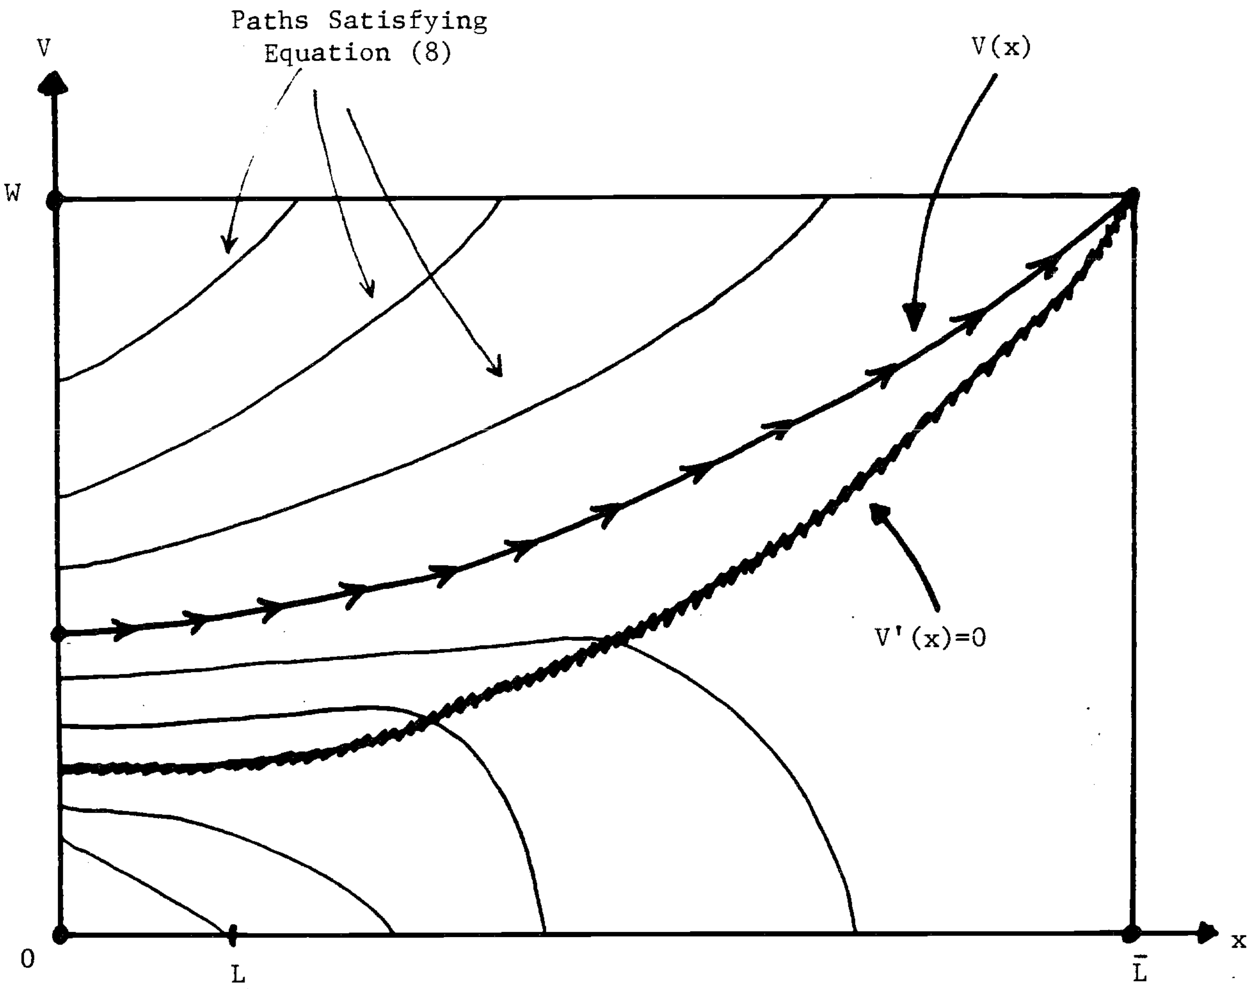

Non-Decreasing Hazard Rate

Figure 2

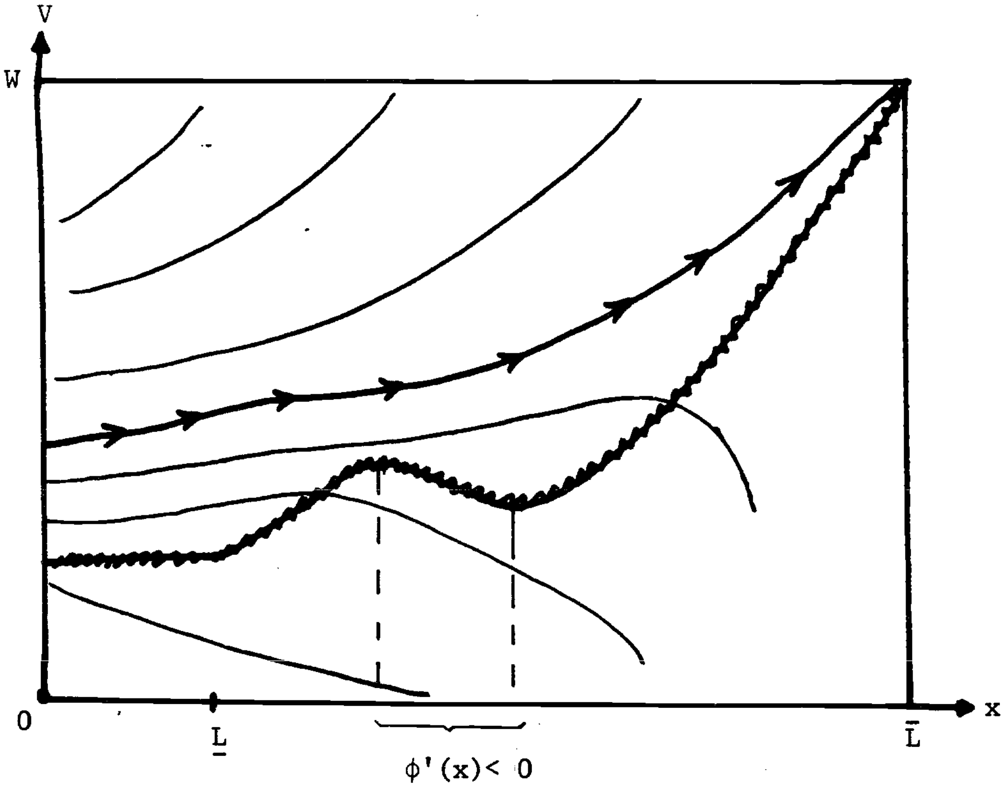

Figure 3a: Hazard Rate Falls

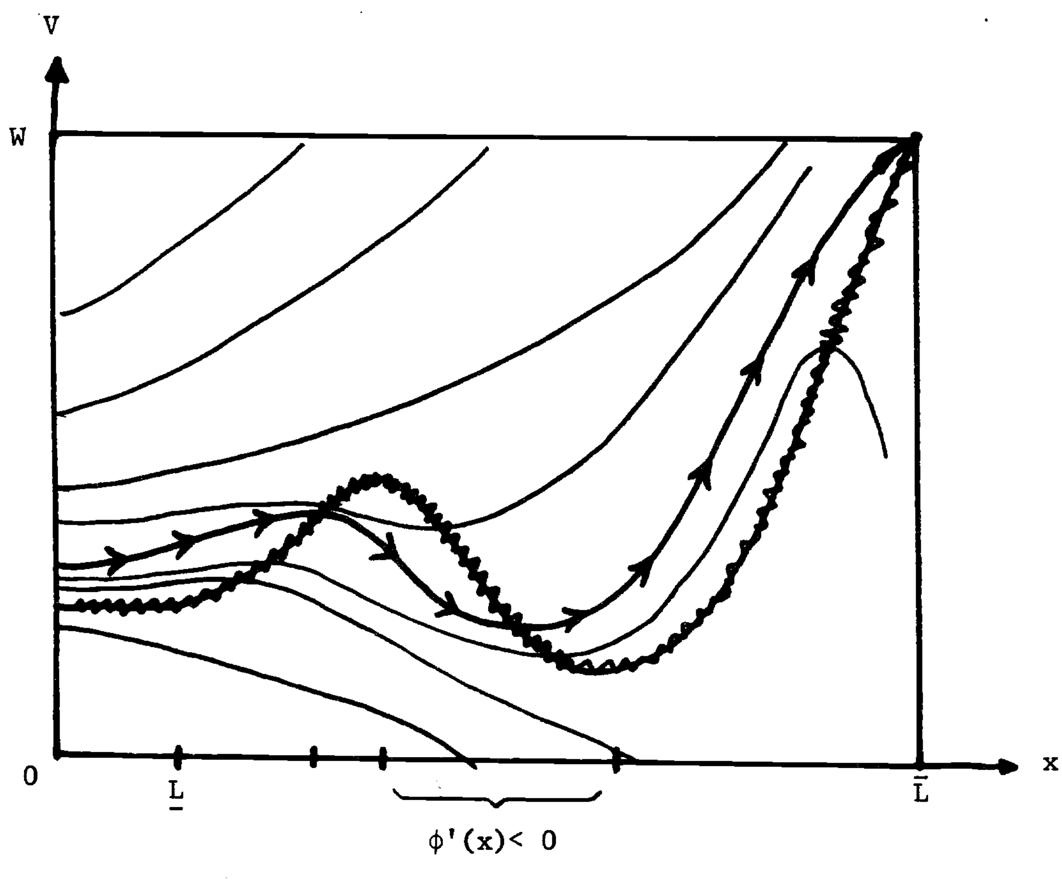

Figure 3b: Hazard Rate and Expenditure Rate Fall

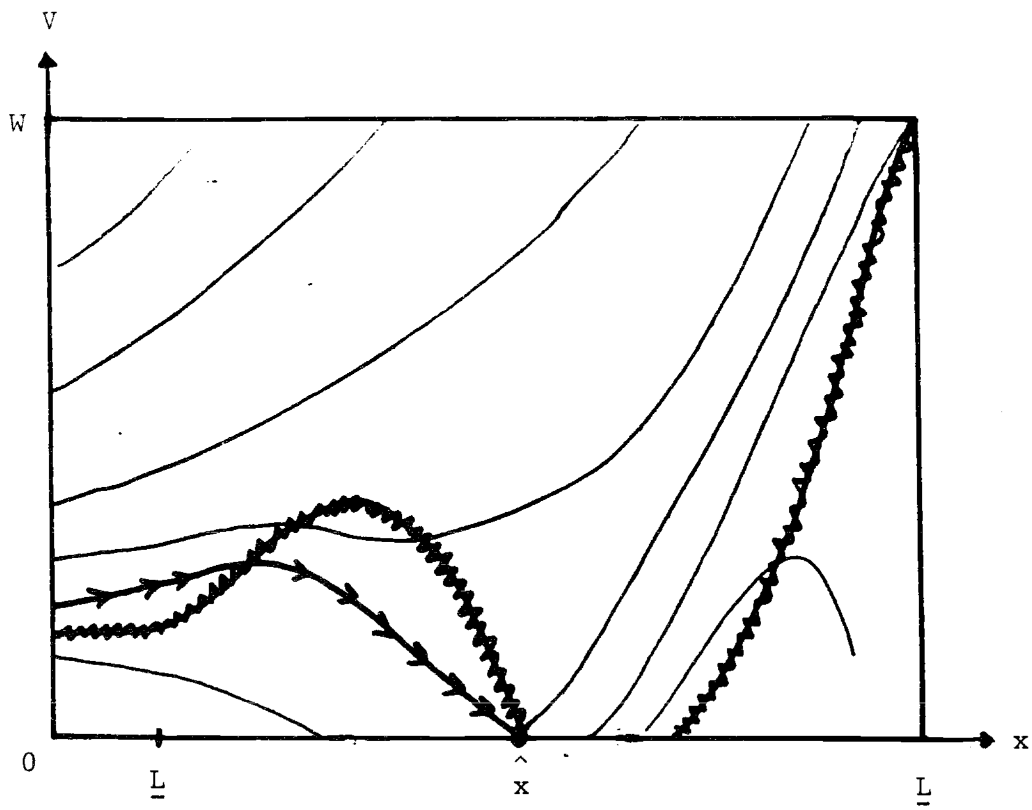

Figure 3c: Abandon Research at $\hat{x}$

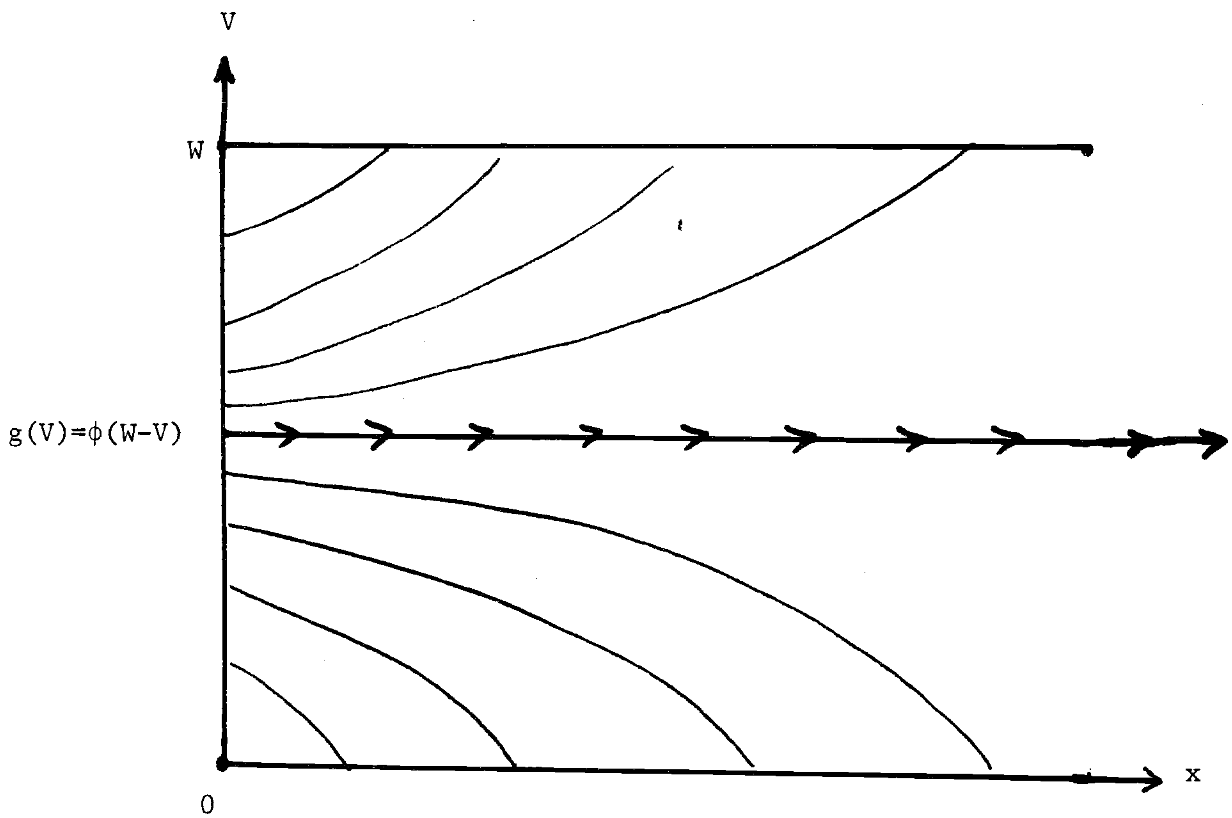

Figure 4: Constant Hazard Rate

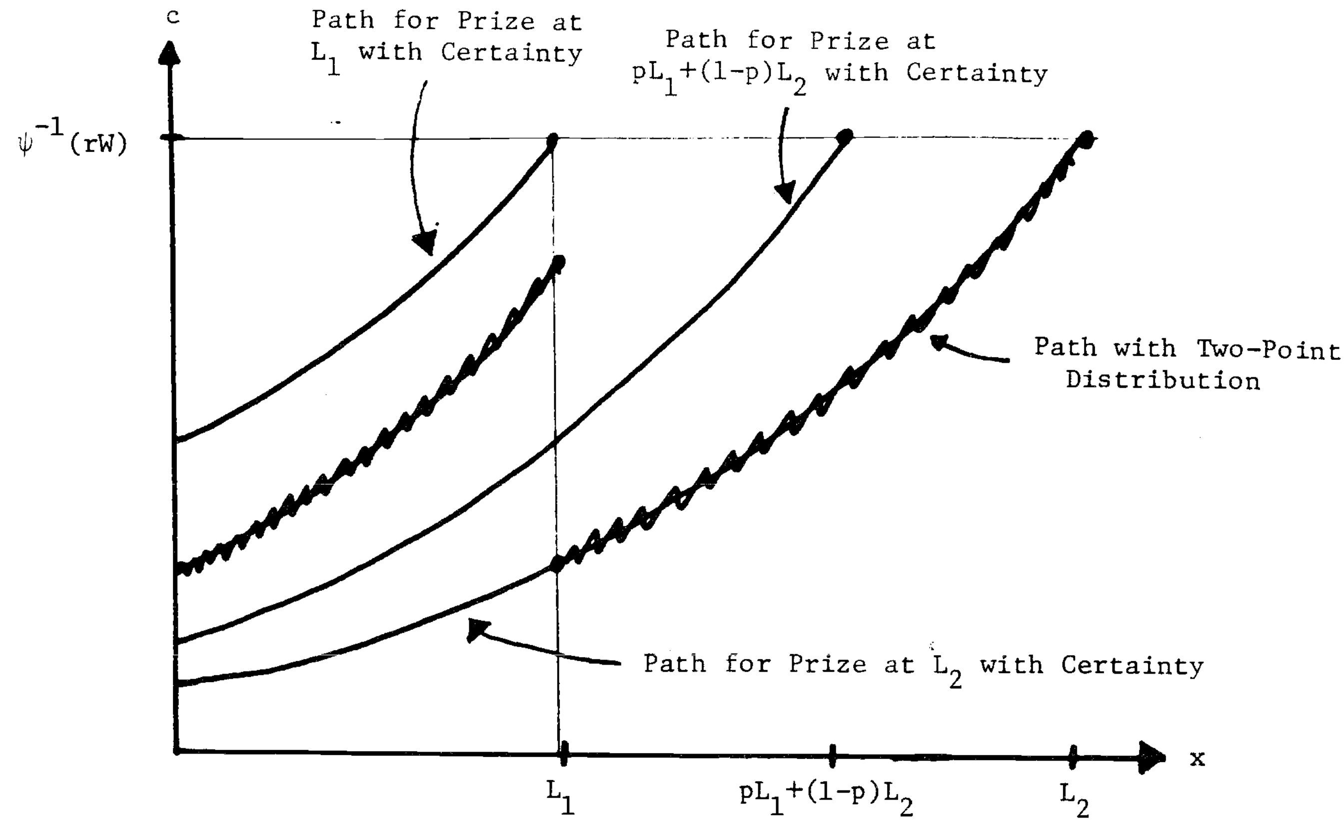

Figure 5: Expenditure Pattern for a Two-Point Distribution

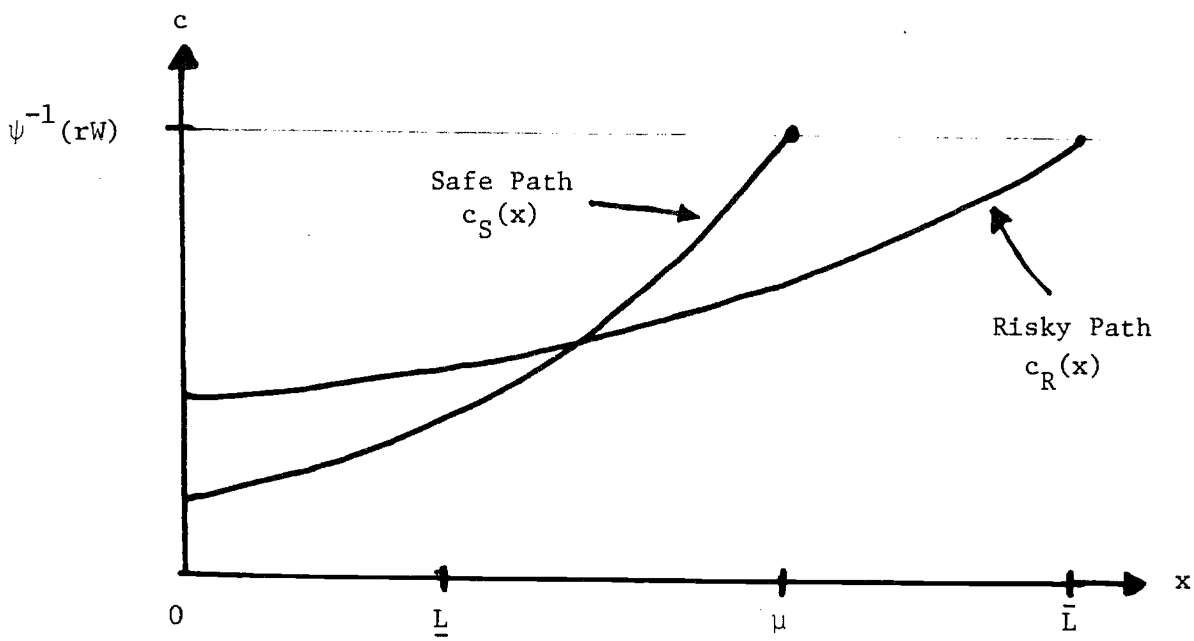

Figure 6: Safe vs. Risky Expenditure Patterns