Rochester Center for
Economic Research

Optimal Management of an R&D Budget

Dutta, Prajit K.

Working Paper No. 357
August 1993

University of
Rochester

Optimal Management of an R&D Budget

Prajit K. Dutta

Rochester Center for Economic Research
Working Paper No. 357

Optimal Management of an R&D Budget*

by

Prajit K. Dutta
Department of Economics
Columbia University
New York, NY 10027

August 1991
Revised: May 1993

Abstract: This paper characterizes the optimal allocation of a budget between several stages of an R&D project. When completion of intermediate stages is profitable, the optimal allocation involves greater expenditures in early stages of the project. Increasing the budget size does not necessarily increase expenditures in every stage (even if the innovation production function is subject to decreasing returns) but does imply that the budget remaining unspent after every stage is greater. A simple necessary and sufficient condition is given for the optimality of "bold play"– the spending of all remaining budget on the current stage. If profits are realized only after all stages are successfully completed, decreasing returns in the innovation production function implies that the optimal strategy is to spread the budget evenly between the stages of the R&D project.
JEL Classification Numbers 020, 621
Keywords Sequential R&D, Dynamics, Budget

* This is a considerably revised version of an earlier manuscript titled "Budget Allocation Under Sequential Uncertainty." I have benefitted from helpful conversations with Peter Linhart and Roy Radner as well the comments of seminar participants at Columbia, Cornell, NYU, Indian Statistical Institute, New Delhi and the NSF–NBER Decentralization Conference held at the University of Rochester. I thank Andrew Schotter for asking me the initial question that led eventually to this paper.

1. Introduction

Many R&D projects are sequential in character; there are several steps to a project and a natural progression which requires completion of the n–th step before inception of the n+1–st. In some cases the project is successfully completed, and payoffs are realized, only if each of these intermediate steps is successfully finished. In some other instances, completion of an intermediate step might itself be profitable. The issue I am interested in is, how much of resources ought to be allocated at each stage of an R&D project?

In order to analyze this question I assume that the total amount of resources available for allocation between the different stages is fixed a priori, i.e., that the real decision is: **how should the project's budget be distributed among its various parts?** There is, to my mind, at least three reasons why we should be interested in the budget allocation problem (as contrasted with an allocation process in which the financing for each atage is decided independently). Casual empiricism suggests that in many instances R&D financing is indeed project–specific. This seems especially so when the actual implementation of an R&D project is in the hands of an **agent** (a research team or an outside laboratory e.g.). In such cases specifying a total budget may be a response to the attendant agency problems (more on this in Section 7). Finally, academic research has almost exclusively taken the stage by stage perspective on R&D financing (for example, see Grossman and Shapiro (1986), Harris and Vickers (1986) and Kamien and Schwartz (1982)). The budget perspective taken in this paper can therefore be viewed, at the very least, as a robustness check on the available literature.

There are only a few papers that develop the intertemporal nature of R&D activity. One exeception is Grossman and Shapiro (1986) which analyses expenditure allocation in a sequential R&D problem. One of the models I study is similar to the one employed in Section 5 of that paper; the important difference is that they did **not** impose any **aggregate budget constraint.**[1] In turn, Grossman and Shapiro were a generalization of earlier work by

Lucas (1971). Mention should also be made of two other papers which model the sequential nature of R&D explicitly; Gallini and Kotowitz (1985) and Granot and Zuckerman (1991). In these papers, at each stage of the R&D project, a decision–maker chooses one of a finite set of available processes, i.e., the central question there is the **sequencing** of stages in an R&D project; my focus is on the intensity of R&D effort or expenditure at each stage of a given sequence of stages.[2] Finally, there is a (non–R&D) literature in operations research which examines the sequential budget allocation problem. The prototypical problem in that literature has been (colorfully) nicknamed the "bomber problem" since it analyzes the optimal rationing scheme of a gunner faced with a sequence of enemy bombers and a fixed cache of ammunition – see Simmons and Yao (1990) and Shepp, Simmons and Yao (1990) (as well as Ross (1983, Chapter 1) for a non–bomber example). (These papers are discussed in greater detail in Section 7.)

Like all of the above papers, I adopt the framework of the decision–theoretic literature on R&D and treat dynamic R&D investment as the outcome of an optimization problem faced by a single firm. The analysis applies immediately to a monopolist undertaking R&D or to a competitive firm that believes its decisions leave its rivals' R&D expenditures unchanged or to a research group or manager whose payoffs are directly determined by the success of the R&D project it is charged with implementing. In Section 7, I will return to a further discussion of these interpretations as well as the significance of my results for game–theoretic and information–theoretic analyses of R&D. I will not distinguish, for now, between the different interpretations and in the sequel I will refer to the decision–maker variously as the firm or R&D manager.

The structure of the allocation problem I study is as follows: in each period an allocation is made (out of the remaining budget) and the size of this allocation determines the probability of successfully completing the current stage (on that attempt). Naturally, the higher is the allocation the more likely is success but the smaller consequently is the size of the budget remaining. I study two polar specifications of the number of attempts that can be made at completing each stage: the **"single attempt"** model in which the first failure at any stage prevents the project from going any further and the **"infinite attempt"** model in which any number of attempts can be made till a stage is (finally) cleared.

I start in Section 3 with the problem in which intermediate steps are valuable; completing a stage yields a (stage–dependent) profit. This seems a good description of a development project and I will call this the **flow payoff** model. In this model, when only single attempts are permissible, the optimal time–path of R&D expenditures is actually **decreasing**, i.e. the earlier the stage, the greater is the allocation to that stage and this is true **regardless of the properties of the stage payoffs**. With infinite mistakes, the same conclusion holds provided the stage payoffs are geometrically increasing or the allocations are stage–specific.

I investigate next in Section 4 the comparative dynamic properties of the optimal allocation and in particular the consequences of increasing the initial budget size. An example demonstrates that a higher budget does not necessarily imply a higher allocation at every stage (and this even if the probabilty of success function is subject to decreasing returns). However, the **budget remaining** after each stage is greater, the bigger the initial budget. These results apply to both the single and infinite attempt versions of the problem.

Of some interest is the possible optimality of exclusive budgeting, i.e. the spending of all remaining budget on the current stage – this behavior has been termed **"bold play"** in the gambling literature (see Dubins and Savage (1965)). I give a simple **necessary and sufficient condition for the optimality of bold play** for an R&D manager: the **marginal** probabilty of success to zero allocation should be **finite**. This set of issues, still within the flow payoff model, is discussed in Section 5.

Section 6 analyzes the **terminal payoffs** model in which profits are realized only after all stages have been successfully concluded. This model is arguably a better description of a pure research problem. I demonstrate that under a log–concavity assumption on the probabilty of success function, (satisfied in particular if the function is subject to decreasing returns) the optimal expenditure pattern is to **spread the budget evenly** among all stages of an R&D project.

Section 7 presents some extensions of the model and a discussion of the literature while Section 8 concludes.

2. The Model

Let $y > 0$ denote a given budget. Informally, a feasible allocation is a distribution of the budget over the different stages of the project such that the total allocation is no greater than y. The precise definition of a feasible allocation is a little bit more complicated since I allow multiple attempts at completing any one stage.

Let $n = 1,..N$ denote the stages of the R&D project (where $N \leq \infty$) and let $t = 1, 2, ...$ denote time–periods. By the beginning of time T, evidently $n-1 \leq T-1$ stages are complete and $y_T \equiv \sum_{t=1}^{T-1} x_t$ is the remaining budget, where x_t denotes the allocation in period t. The allocation in period T is made contingent on the current stage, i.e. n, and the remaining budget, i.e. y_T, and is denoted $\alpha(y_T,n)$.[3] Of course, it is further required that $0 \leq \alpha(y_T,n) \leq y_T$. An **R&D budget strategy** for the firm is a sequence of such allocation functions $\underset{\sim}{\alpha} \equiv [\alpha(.,n)\colon n = 1,..N]$. The allocation in period T suffices to complete the n–th stage with **probability** $\rho(\alpha(y_T,n))$ and with the remaining probabilty, the project remains at stage n. Suppose also that success in different periods are independent events.

I examine two different specifications of payoffs:

Flow Payoffs Completing stage n yields a profit of $\gamma_n > 0$. This profit may be viewed either as a once–for–all return or the present discounted value of all future returns generated by the completion of this stage. Clearly, the expected return in period T, conditional on being at (y_T,n) and the consequent expenditure, is $\rho(\alpha(y_T,n))\gamma_n$.

Furthermore, any budget strategy $\underset{\sim}{\alpha}$ and the success probability function ρ define a period–T distribution over the set of remaining budgets and the number of completed stages[4]– and hence yield an expected profit for period T which I call $\Gamma_T(y;\underset{\sim}{\alpha})$. The expected lifetime payoffs to the R&D allocation is therefore

$$R(y;\underset{\sim}{\alpha}) = \sum_{t=1}^{\infty} \delta^{t-1} \, \Gamma_t(y;\underset{\sim}{\alpha}) \tag{1.1}$$

The objective for the firm is to maximize the expected lifetime payoffs by the choice of budget strategy. In order to fix ideas, let me report the precise form of (1.1) in the two polar cases of single and infinite admissible attempts. When the first failure terminates the R&D project, the expected lifetime allocation becomes

$$\begin{aligned} R(y;\underset{\sim}{\alpha}) &= \rho(\alpha(y,1))\gamma_1 + \delta[\rho(\alpha(y,1)) \times \rho(\alpha(y_2,2))]\gamma_2 + \ldots\delta^n \prod_{\zeta=1}^{n} \rho(\alpha(y_\zeta,\zeta)\gamma_n + \ldots \\ &= \sum_{n=1}^{N} \delta^{n-1}\gamma_n \prod_{\zeta=1}^{n} \rho(x_\zeta) \end{aligned}$$

where x_ζ is simply $\alpha(y_\zeta,\zeta)$. On the other hand, if infinite attempts are permissible then the expected lifetime payoffs become

$$R(y;\underset{\sim}{\alpha}) = \rho(\alpha(y,1))\gamma_1 + \delta[\rho(\alpha(y,1))\times\rho(\alpha(y_2,2))]\gamma_2 + \delta[(1-\rho(\alpha(y,1)))\times\rho(\alpha(y_2,1))]\gamma_1 + \ldots$$

A second formulation of returns to R&D budget allocation is:

Terminal Payoffs Here returns only accrue after all stages are complete (and so this formulation is appropriate only when $N < \infty$). Let this terminal payoff be denoted W. It is clear that any budget strategy $\underset{\sim}{\alpha}$ generates a probability that the entire project is in fact completed at period t; denote this probabilty μ_t. Then, the expected lifetime profits are

$$r(y;\underset{\sim}{\alpha}) = \sum_{t=1}^{\infty} \delta^{t-1} \, \mu_t W \tag{1.2}$$

As before, in order to fix ideas consider the case of a single permissible attempt at each stage. In that setting,

$$r(y;\underset{\sim}{\alpha}) = \delta^{N-1}[\prod_{\zeta=1}^{N} \rho(x_\zeta)]W$$

The following two, mild, **assumptions** will be maintained on the probability of success function ρ:

(A1) ρ is a continuous function

(A2) ρ is weakly monotonic, $x' \geq x$ implies $\rho(x') \geq \rho(x)$.

Remark: (A1) is only required in order to ensure that the budget allocation problem has a solution. The monotonicity assumption, (A2), is in fact an assumption that can be made without loss of generality; if ρ were actually declining over some allocations, then, in an optimal solution, expenditures would never be made from such decreasing segments.

In Section 7, I discuss extensions of the analysis to the case in which the probabilty of success at any stage may depend on the stage itself (in addition to the allocation made at that stage). Also, at various points of the discussion, I will mention the extent to which the results remain unchanged if the framework above is enlarged to admit a scrap–value payment which is made on the budget remaining unused when the project is either completed or terminated. Section 7 also contains a discussion of a number of other directions for future research.

3. The Time–Path of Optimal Allocations

In this section I derive the time–path of optimal R&D expenditures when completing intermediate steps of the project yield flow payoffs. This is the appropriate framework, for instance, for a sequential development project, each step of which is the further improvement of a basic technology or the further increment of quality of product.

It is evident that the flow payoff problem can be cast in the language of dynamic programming with the remaining budget and the number of the current stage as the state variable and the immediate allocation as the action or control variable. In order to ensure that expected lifetime payoffs are well–defined (in the potentially troublesome case where $N = \infty$), I assume that there is an upper bound M such that for all $n \geq 1$, $\sum_{t=1}^{\infty} \delta^{t-1} \bar{\gamma}_t(n) < M$, where $\bar{\gamma}_t(n) \equiv \max(\gamma_{n+1},..\gamma_{n+t})$. Armed with this assumption, a direct appeal to standard results in dynamic programming implies that there is a continuous value function, denoted $V(y,n)$, and an optimal budget allocation function $\alpha^*(y,n)$ (see, for example, Stokey et. al. (1989, Theorem 4.14) for details).[5]

Consider, to begin with, the allocation problem when the first failure terminates the R&D project. In this case, the actual allocations are as follows: $x_1^* = \alpha^*(y,1)$, the remaining budget y_t is defined inductively as $y_{t+1} = y_t - \alpha^*(y_t,t)$ and the allocation at the t+1 stage (and time–period) is $x_{t+1}^* = \alpha^*(y_{t+1},t+1)$. Of course, a stage is reached only if all previous stages have been successfully completed. In principle, there may be several optimal budget allocation policies $\alpha^*(y,n)$, and I shall call the time–allocation implied by any one of them an **optimal allocation**. The following result characterizes the time–path of optimal R&D allocations:

Proposition 1 **Under (A1) and (A2), there is at least one optimal allocation (x_t^*, $t \geq 1$) such that**

$$x_t^* \geq x_{t+1}^*, \quad t \geq 1. \tag{3.1}$$

If the probability of success function is strictly monotonic, then all optimal allocations satisfy (3.1).

Proof: Suppose that $\alpha^*(y,n)$ is an optimal budget strategy with an associated allocation x_t^*, $t \geq 1$. Consider any stage n (equivalently, time–period $\zeta = n$) and further, consider an alternative strategy, say $\hat{\alpha}$, whose allocation is identical to that of α^* except that the period ζ and $\zeta+1$ allocations are switched; $\hat{x}_\zeta = x^*_{\zeta+1}$, $\hat{x}_{\zeta+1} = x^*_\zeta$ and $\hat{x}_t = x^*_t$ for all $t \neq \zeta, \zeta+1$, $\zeta < N$. Evidently, this is a feasible budget strategy. Further, the probability of succeeding in all of the first t stages remains completely unchanged provided

$t \neq \zeta$ and when $t = \zeta$, the probabilities of success differ by $[\rho(x^*_\zeta) - \rho(x^*_{\zeta+1})] \prod_{t=1}^{\zeta-1} \rho(x^*_t)$. It follows then that

$$R(y;\alpha^*) - R(y;\hat{\alpha}) = \delta^\zeta[\rho(x^*_\zeta) - \rho(x^*_{\zeta+1})][\prod_{t=1}^{\zeta-1} \rho(x^*_t)]\gamma_\zeta \tag{3.2}$$

From (3.2) it is evident that the optimality of α^* implies $\rho(x^*_\zeta) \geq \rho(x^*_{\zeta+1})$, for all $\zeta < N$. In the light of the weak monotonicity assumption (A2), this says that $x^*_\zeta < x^*_{\zeta+1}$ is only possible if the two allocations lie in an interval over which the function ρ is constant. But in that instance, we can amend α^* by requiring that it use the smallest allocation consistent with that constant probability of success as the allocation at stage $\zeta + 1$. Furthermore, it is clear that by adjusting the allocations in such a manner at all time–periods where α^* takes values in a "flat" of the probabilty function, we have an optimal R&D allocation strategy which satisfies (3.1).

If ρ is strictly increasing (3.2) evidently implies (3.1). □

Remark Suppose the stage payoffs γ_n depend on the size of the allocation x_n as well. The argument, and hence the result, is unchanged if γ_n is a non–decreasing function. As a second generalization, consider the possibility that if the R&D project is terminated (upon failure at some stage), then a termination or scrap–value payment is made, say S_n, which is contingent on the number of completed stages. It is straightforward to check that the above argument remains completely unchanged.[6]

The reason why an R&D manager will allocate greater funds to earlier stages of the project is quite straightforward; conditional on successful completion of the t–th stage, the expenditures in stages t and t+1 have a symmetric effect on all subsequent payoffs. However, the funds spent in stage t additionally determine the likelihood of successfully completing the t–th stage.

The conclusion that more of the budget is spent in earlier stages is not an artifact of the requirement that each stage can be attempted only once. To see this, I establish an analogous property in the model where the firm gets many attempts to complete each stage and in fact gets as many attempts as it needs. Expositionally, an immediate problem is that the actual allocation of the R&D budget is now a stochastic process and hence the notion of a time–pattern of allocation is ambiguous. Moreover, the optimal budget strategies $\alpha^*(y,n)$ will not, in general, satisfy convenient monotonicity properties; for example of the form the smaller the number of completed stages, the higher (or lower) is the allocation out of any remaining budget.[7] There are however two important formulations of the infinite attempt problem which are tractable and whose conclusions mirror those under single attempts.

In the first formulation, following Dasgupta and Stiglitz (1980) or Kamien and Schwartz (1982), I make the simplifying assumption that the **total** allocation for any stage is all that the R&D manger can choose. In other words, at the beginning of stage n, the R&D manager picks an allocation x_n which is targetted towards completing that stage. The number of attempts, T, it takes to actually complete the stage is given by a distribution function $F(.;x_n)$ (with an associated density, say $f(.;x_n)$). A higher allocation is (stochastically) more likely to complete a stage sooner, i.e. the analog of (A2) is

(A2') $x' \geq x$ implies that the distribution $F(.;x')$ first–order stochastically dominates the distribution $F(.;x)$.

Write x_n^* for the optimal allocation at stage n, i.e. $x_n^* = \alpha^*(y_n^*,n)$ and note that

$$V(y,n) = [\gamma_n + V(y-x_n^*,n+1)]\,[\Sigma_{T=1}^{\infty}\ \delta^T f(T;x_n^*)] \qquad (3.3)$$

By virtue of (A2'), the expected discounted time to success, $\Sigma_{T=1}^{\infty}\ \delta^T f(T;x)$, is a non–decreasing function of the allocation x. To conserve exposition I will refer to the expected discounted time to success as $\tau(x)$ from this point on. It is now clear from (3.3) that the structure of the problem is identical to the single attempt case with the expected discounted time to success playing here the role that the probability of success function played in the previous formulation. Identical arguments as in Proposition 1 yield:

Proposition 2 **Under (A1) and (A2'), optimal budgetary allocations are non–increasing in the number of completed stages, i.e.** $x_n^* \geq x_{n+1}^*$.

A second tractable formulation allows the decision–maker to adjust his expenditures within any one stage but requires a geometric structure of the payoffs:

(A3) The stage payoffs increase geometrically; $\gamma_{n+1} = \beta\gamma_n$, where $\beta \geq 1$.[8] Further, there are an infinite number of stages, i.e. $N = \infty$.

In this formulation, the agent's decision rules have the following stationarity property:

Proposition 3 **Under (A1)–(A3), optimal budgetary strategies are independent of the number of completed stages, i.e. are of the form** $\alpha^*(y,n) = \alpha^*(y,n+1)$, **for all y, n < N. Furthermore, the value function is multiplicative in budget size and the number of the current stage,** $V(y,n) = \beta^{n-1}V(y,1)$.

Proof: In Appendix B. □

In the light of Proposition 3 it is clear that the optimal R&D allocation in any time–period t is a constant, x_t^*. Of course, which stage of the project this constant allocation is being employed for depends on the past history of successes and failures. The following result characterizes the time–path of optimal R&D allocations:

Proposition 4 **Under (A1)–(A3), there is at least one optimal allocation which assigns decreasing expenditures over time, i.e.** $x_t^* \geq x_{t+1}^*$. **If the probabilty of success function is strictly increasing, then all optimal allocations have this feature.**

Proof: The optimality equation, evaluated at period t, yields

$$V(y_t,n) = \rho(x_t)\gamma_n + \delta\left[\{1-\rho(x_t)\}\rho(x_{t+1})\gamma_n + \rho(x_t)\rho(x_{t+1})\gamma_{n+1}\right]$$

$$+ \delta^2\Big[\rho(x_t)\rho(x_{t+1})V(y_{t+2},n+2) \quad + \quad \{[1-\rho(x_t)]\rho(x_{t+1})+[1-\rho(x_{t+1})]\rho(x_t)\}V(y_{t+2},n+1)$$

$$+ [1-\rho(x_t)][1-\rho(x_{t+1})]V(y_{t+2},n)\Big] \tag{3.4}$$

Consider, as in the proof of Proposition 1, switching the allocations in period t and t+1; so, in the alternative strategy $\hat{\alpha}$, $\hat{x}_\zeta = x^*_\zeta$, $\zeta \neq$ t, t+1 and $\hat{x}_t = x^*_{t+1}$, $\hat{x}_{t+1} = x^*_t$. This leaves y_{t+2} unchanged, as also the multiplicative probabilities $\rho(x_t)\rho(x_{t+1})$ and $[1-\rho(x_t)][1-\rho(x_{t+1})]$. Using that information it is evident that all terms in (3.4), except the first two, are in fact identical. In other words, the difference in expected payoffs, starting at period t, between the optimal strategy and the switched strategy is simply $\gamma_n(1-\delta)[\rho(x_t) - \rho(x_{t+1})]\gamma_n$. From that it follows that $\rho(x_t) \geq \rho(x_{t+1})$. Having established the time–pattern of success probabilties, the remaing arguments that establish the time–pattern of optimal allocations are identical to those used in the analogous part of the proof of Proposition 1. □

4. The Effects of Changing The Size of the Research Budget

Suppose the initial budget y is increased. In this section I examine the effect of a change in budget size on its allocation; in particular, I ask (when) is it the case that increasing y, a) increases the optimal allocation $\alpha^*(y,n)$ and/or b) increases the size of unspent budget $y - \alpha^*(y,n)$ (in each case, for all n). A corollary of the latter is evidently that the entire time–path of optimal unused budget, $[y^*_t,\ t \geq 1]$, is shifted up with an increase in the size of the initial budget. In turn, if both a) and b) hold, then the entire time–path of budget usage, $[x^*_t,\ t \geq 1]$, shifts up with an increase in the initial budget.

The critical factor turns out to be whether or not there are increasing returns in the probabilty of success function ρ. Now the probabilty of success function is, in my model, the analog of an "innovation production function." Whether or not it is subject to increasing returns was the subject of considerable empirical research in the 1960's (see Kamien and Schwartz (1982, Chapter 3) for a lucid review). That literature suggested that in some industries there is indeed evidence of increasing returns to R&D production (at least initially) although many others appear to exhibit decreasing returns throughout. I will show that in the presence of decreasing returns to the probability of success function, **unspent budget is increasing in budget size.** However, even with such decreasing returns, **budget expenditures need not be increasing in budget size.**

I start with the second question, present a counter–example when ρ has increasing returns and then prove a monotonicity result for concave ρ. Note that if the probability of success function is convex, at least over some allocations, then the optimal usage may involve immediate exhaustion of the entire budget. However, if increasing returns are strong only for sufficiently large allocations, then a smaller budget need not be immediately exhausted in any optimal scheme. Hence, with non–convexities in the probability of success function, unused budget may not be larger from a bigger initial budget. Example 1 makes precise this intuition.

Example 1. **The optimal budget allocations are such that the unspent budget is smaller when the initial budget is in fact larger.**

Details: Let ρ be defined on the domain [0, 1] by the requirements that i) it has a strictly concave segment on [0, 1/2] followed by a strictly convex segment on [1/2, 1]. Further, ii) ρ is strictly increasing and symmetric about $x = 1/2$, iii) $\rho(0) = 0$ (and $\rho(1) = 1$) and iv) ρ is differentiable with $\rho'(1/2) = 0$. (All of these conditions are satisfied, for example, by $\rho(x) = 1/2 + 4(x - 1/2)^3$). Further, for simplicity, consider the single attempt case when $N = 2$, $\delta = 1$ and the stage payoffs are a constant (say 1). Under the above restrictions, ρ looks as in Figure 1:

(Figure 1)

Suppose the initial budget is 1. Then, $R = \rho(x) + \rho(x)\rho(1-x) = \rho(x)[2 - \rho(x)]$, where the last equality utilized the fact that ρ is symmetric about $x = 1/2$. Since $\rho(x) \in$

[0,1] clearly the optimal first period allocation $x^* = 1$, i.e. the budget is completely exhausted. Suppose instead that the initial budget is 1/2. I claim that the optimal allocation leaves some budget unspent; to see this note that if the optimal allocation were exhaustion of budget, then marginal payoff at that utilization must be non–negative, i.e. $\frac{d}{dx}R|_{x = 1/2} \geq 0$. However, $\frac{d}{dx}R|_{x = 1/2} = -\rho(1/2)\rho'(0) = -(1/2)\rho'(0) < 0$. □

The driving force behind Example 1 is of course the increasing returns to probabilty of success, for allocations in [1/2, 1]. If the probabilty of success function is subject to decreasing returns, then the larger the budget the greater the amount left unspent after any period's or any stage's allocation. Note, incidentally, that whenever I refer to the multiple attempt formulation I speak of the two frameworks covered by (either) (A2') or (A3).

Proposition 5 **Consider either the single attempt or the multiple attempt cases. Suppose, additionally, that ρ (respectively, τ) is strictly concave.[9] Then in any optimal budget strategy α^*, $y - \alpha^*(y,n)$ is non–decreasing in y, for all n.**

Proof: I present the argument only for the single attempt case. A straightforward adaptation of the argument applies in the multiple attempt problem. Let α^* be any optimal budget strategy and consider two initial budgets y and y', with $y > y'$. Let $z = y - \alpha^*(y,n)$ and $z' = y' - \alpha^*(y',n)$. For simplicity of notation, in what follows I drop the argument n from the value functions and write V(y) instead of V(y,n) throughout. Suppose, in contradiction to the result claimed, $z < z'$. From the optimality equation it follows that

$$\rho(y - z)\gamma_n + \delta\rho(y - z)V(z) \geq \rho(y - z')\gamma_n + \delta\rho(y - z')V(z')$$
$$\rho(y' - z')\gamma_n + \delta\rho(y' - z')V(z') \geq \rho(y' - z)\gamma_n + \delta\rho(y' - z)V(z)$$

Combinining the two inequalities yields

$$[\rho(y - z) - \rho(y' - z)][\gamma_n + \delta V(z)] \geq [\rho(y - z') - \rho(y' - z'][\gamma_n + \delta V(z')]$$

The inequality above, together with the fact that V(.,n) is strictly increasing, implies that

$$[\rho(y - z) - \rho(y' - z)] > [\rho(y - z') - \rho(y' - z'] \tag{4.1}$$

(4.1) yields a contradiction in light of the strict concavity of ρ. □

As noted above, an obvious corollary to the above proposition is the following: the time–path of optimal allocations shifts monotonically up when the initial budget size is increased; $y_t \geq y'_t$, $t \geq 1$.

The presence of decreasing returns in ρ does not however guarantee that an increase in the budget necessarily increases the size of allocation out of it; i.e. it does not guarantee that $\alpha^*(y,n)$ is non–decreasing in y. It turns out that such monotonicity in usage is determined by whether or not the value function has decreasing returns. Even if the probability of success function is concave, the value function is not in general concave since the objective function R(·) contains terms involving the **product** of probabilities of success and hence may not be concave. In the event that the value function is convex, an increase in the initial budget may induce the R&D manager to **reduce** his current allocation in order to exploit the increasing returns present in the value function.

I now present a three–stage example whose optimal stage one budget strategy has the following feature: for "small" y all of the budget is allocated in the first stage itself. At a critical budget size $\hat{y}$, the R&D manager is indifferent between spending all of his budget and spending only a fraction of it, $x_1(\hat{y}) < \hat{y}$. Budgets in excess of $\hat{y}$ are allocated in the first stage according to the (continuous) function $x_1(.)$. Hence, the optimal budget strategy is **discontinuous** and higher budgets can lead to **smaller allocations**.

Example 2 The optimal budget allocation α^* is not non–decreasing even though ρ is strictly concave.

Details: Let $\rho(x) = \dfrac{x}{x+1}$, a strictly increasing and strictly concave function.

Consider the single attempt problem and let there be three stages. Further, to facilitate computation, let $\gamma_1 = \gamma_3 = 1$ and $\gamma_2 = 0$ and suppose that $\delta = 1$. When two stages are left, the optimization problem out of a remaining budget of size y is, $\underset{x \in [0,y]}{\text{Max}} \rho(x)\rho(y-x)$. From the strict concavity of ρ it follows that the optimal choice at that stage is to spread the budget evenly between the two stages, i.e. the optimal choice is $x^* = y/2$. Notice that, as a consequence, $V(y,2) = (\frac{y}{y+2})^2$, a convex function for $y > 1$.

The three stage optimization problem is hence given by

$$\underset{x \in [0,y]}{\text{Max}} \frac{x}{x+1} [1 + (\frac{y-x}{y-x+2})^2] \tag{4.2}$$

In Appendix A I demonstrate that the objective function in (4.2) has, for every y, one local maximum at $x_1 \in (0,y)$ and a local minimum at $x_2 \in (x_1,y)$. In other words, the function looks as follows

(Figure 2)

Evidently, the solution to the optimization problem is either x_1 or y. I further show in the appendix that the optimal budget strategy is (essentially) given by: there is $\hat{y} > 0$, such that for all $y < \hat{y}$, the optimal allocation is y, whereas for $y > \hat{y}$, the optimal allocation is x_1 and at $\hat{y}$ the budget manager is indifferent between these two options. Evidently then, the optimal allocation "jumps down" at $\hat{y}$. □

If the value functions happen to be concave, then monotonicity does obtain in the usage of the budget:

Proposition 6 Consider either the single attempt or the multiple attempt case. Suppose, additionally that V(.,n) is strictly concave for all n. Then, in any optimal budget strategy α^*, $\alpha^*(y,n)$ is non–decreasing in y, for all n.**

Proof: I prove the single attempt case since the multiple attempt problem is again a minor variant. Further, the single attempt proof is itself very similar to the proof in

Proposition 5 and is therefore only sketched. Let y > y' and denote $\alpha^*(y,n)$ as x (respectively, $\alpha^*(y',n)$ as x'). Suppose, in contradiction to the result claimed, x < x'. Since x is a feasible allocation from y' and x' is similarly feasible from y, the optimality equation yields

$$\rho(x)\gamma_n + \delta\rho(x)V(y-x) \geq \rho(x')\gamma_n + \delta\rho(x')V(y-x')$$
$$\rho(x')\gamma_n + \delta\rho(x')V(y'-x') \geq \rho(x)\gamma_n + \delta\rho(x)V(y'-x)$$

Combinining the two inequalities yields

$$\rho(x)[V(y-x) - V(y'-x)] \geq \rho(x')[V(y-x') - V(y'-x')]$$

The inequality above yields a contradiction given the strict concavity of the value function. □

5. Exhausting the Budget

Will it ever be optimal to spend all of the remaining budget at one go, whenever there is less than some critical amount, say $\hat{y}$, left? In the mathematical analyses of gambling such behavior is characterized "bold play" (see, e.g., Dubins and Savage (1965)) and there is a substantial literature that investigates conditions under which such behavior is optimal. Borrowing the term from that literature, I say that **bold play is optimal behavior at stage n** if there exists $\hat{y}$ such that $\alpha^*(y,n) = y$ for all $y \leq \hat{y}$.

For simplicity, in this section I confine discussion to the single attempt case. The generalization to the multiple attempts case is immediate. I strengthen the continuity assumption on the probability of success function to differentiability:

(A1)′ ρ is continuously differentiable for $x > 0$. Further, $\lim_{x \downarrow 0} \rho'(x) \equiv \rho'(0)$ exists and is positive.[10]

Proposition 7 **Under (A1)′ and (A2), for bold play to be optimal at any stage $n < N$ it must be the case that $\rho'(0) < \infty$. Further, if $\rho(0) = 0$, then $\rho'(0) < \infty$ is also sufficient**

for bold play to be optimal at every stage of the R&D project.

Proof: **Necessity:** Consider any stage $n < N$, and suppose that there is $\hat{y} > 0$ such that the optimal allocation starting at $\hat{y}$ is immediate exhaustion of the budget. In particular, this strategy is optimal among those that allocate $x < \hat{y}$ at the current stage and the remainder, $\hat{y} - x$ in the very next stage. In other words,

$$\frac{d}{dx} \left[\rho(x)\, \gamma_n + \delta\rho(x)\, \rho(\hat{y} - x)\, W_n\right]\Big|_{x=\hat{y}} \geq 0$$

where $W_n = \sum_{t=1}^{N-n} (\delta\rho(0))^{t-1} \gamma_{n+t}$. Now

$$\frac{d}{dx}\left[\rho(x)+\delta\rho(x)\rho(\hat{y} - x)W_n\right]\Big|_{x=\hat{y}} = \rho'(\hat{y})[\gamma_n+\delta\rho(0)W_n] - \delta\rho(\hat{y})\rho'(0)W_n \quad (5.1)$$

From (5.1) it immediately follows that $\rho'(0) < \infty$.

Sufficiency: Suppose $\rho'(0) > 0$. I will divide the proof of sufficiency into two parts, one for finite N and the second for infinite N.

Case 1 – $N < \infty$: At stage N it is evidently optimal to exhaust any remaining budget. Make the induction hypothesis that bold play is, in fact, optimal when we are at stage n+1. Let $\hat{y}$ be a budget size below which it is optimal to exhaust the budget; i.e., $\alpha^*(y,n+1) = y$, for all $y \leq \hat{y}$. Now suppose we are at stage n with a budget $y \leq \hat{y}$. Evidently the lifetime returns from an allocation x at stage n is given by $R(x,y) = \rho(x)[\gamma_n + \delta\rho(y-x)W_n]$. I claim that there is always $\tilde{y} > 0$ such that $R(y,x)$ is increasing over $[0,y)$ whenever $y \leq \tilde{y}$. Suppose to the contrary we have a sequence $y_p \to 0$, and x_p in $[0, y_p]$, with $x_p \to 0$, as $p \to \infty$, such that $\frac{d}{dx} R(y_p,x_p) < 0$. Taking limits (and note that the limit is well–defined since $\rho'(0) < \infty$) we have,

$$\lim_{p\to\infty} \frac{d}{dx} R(y_p,x_p) = \rho'(0)[\gamma_n + \delta\rho(0)W_n] - \delta\rho'(0)\rho(0)W_n = \rho'(0)\gamma_n > 0$$

The last inequality yields a contradiction. In turn that establishes, for stage n, the optimality of bold play for all $y \leq \tilde{y}$.

Case 2 – $N = \infty$: We cannot use a backward induction argument for this case. Therefore, we will need the following lemma, which is, in any case, a result of independent interest:

Lemma **For every stage n, the value function, V(.,n) is a differentiable function (on $\mathbb{R}_+$) and the derivative is given by:**

$$V'(y,n) = \rho'(x^*)\,[\gamma_n + \delta V(y-x,n+1)] \tag{5.2}$$

where x^* is the optimal allocation out of a budget of size y at stage n, i.e., $x^* = \alpha^*(y,n)$.

Proof of lemma: Consider any $y > 0$ with an optimal allocation x (for notational simplicity suppress the *). From a budget $y + \epsilon$, $\epsilon > 0$, a feasible (but possibly inoptimal) allocation is $x + \epsilon$. Hence, $V(y+\epsilon) \geq \rho(x+\epsilon)[\gamma_n + \delta V(y-x,n+1)]$. Consequently,

$$V(y+\epsilon) - V(y) \geq [\rho(x+\epsilon) - \rho(x)][\gamma_n + \delta V(y-x,n+1)] \tag{5.3}$$

Taking limits, as $\epsilon \to 0$, it follows from (5.3) that $V'(y) \geq \rho'(x)[\gamma_n + \delta V(y-x,n+1)]$. In order to establish the opposite inequality, note first that a consequence of the assumption that $\rho(0) = 0$ is that $x > 0$. So take any $\epsilon < x$ and consider a budget of size $y - \epsilon$. A feasible allocation is $x - \epsilon$ and hence $V(y-\epsilon) \geq \rho(x-\epsilon)[\gamma_n + \delta V(y-x,n+1)]$. Consequently,

$$V(y) - V(y-\epsilon) \leq [\rho(x) - \rho(x-\epsilon)][\gamma_n + \delta V(y-x,n+1)] \tag{5.4}$$

From (5.4) it follows that $V'(y) \leq \rho'(x)[\gamma_n + \delta V(y-x,n+1)]$. The lemma is proved. □

We now return to the proof of Case 2. The lifetime returns at stage n from a budget y and allocation x is $R(x,y) \equiv \rho(x)[\gamma_n + \delta V(y-x,n+1)]$. I claim that there is always $\tilde{y} > 0$ such that R(y,x) is increasing over [0,y) whenever $y \leq \tilde{y}$. Suppose to the contrary we have a sequence $y_p \to 0$, and x_p in $[0, y_p]$, with $x_p \to 0$, as $p \to \infty$, such that $\frac{d}{dx} R(y_p,x_p) < 0$, i.e., $\rho'(x_p)[\gamma_n + \delta V(y_p-x_p,n+1)] - \delta\rho(x_p)V'(y_p-x_p,n+1) < 0$. Given the lemma above, $V'(y_p-x_p,n+1) = \rho'(\hat{x}_p)[\gamma_n + \delta V(z_p,n+2)]$, where $\hat{x}_p$ is the optimal allocation out of a budget $y_p - x_p$ at stage n+1 and z_p is the consequent budget remaining at stage n+2. Substituting for the above expression and taking limits we have,

$$\begin{aligned}\lim_{p\to\infty}\frac{d}{dx}R(y_p,x_p) &= \rho'(0)[\gamma_n + \delta V(0,n+1)] - \delta\rho(0)\rho'(0)[\gamma_{n+1} + \delta V(0,n+2)] \\ &= \rho'(0)\gamma_n + \delta\rho(0)\rho'(0)[\gamma_{n+1}+\delta V(0,n+2)] - \delta\rho(0)\rho'(0)[\gamma_{n+1}+\delta V(0,n+2)] \quad (5.5) \\ &= \rho'(0)\gamma_n > 0\end{aligned}$$

In (5.5) I used the fact that $V(0,n+1) = \rho(0)[\gamma_{n+1}+\delta V(0,n+2)]$. The last inequality above yields a contradiction. In turn that establishes, for stage n, the optimality of bold play for all $y \leq \tilde{y}$. The proposition is proved. □

Remark: From the proof above it is clear that the condition $\rho(0) = 0$ is only required for the case $N = \infty$ (it was used in proving the lemma and only used then to argue that an optimal allocation must be positive). It is my conjecture that the lemma, and therefore Proposition 7, is true even without this condition.

If the success probability function ρ is **convex** then bold play is optimal regardless of the size of the remaining budget. Of course, since it only takes values in [0,1], ρ cannot be convex throughout but suppose it is convex on an initial segment. (Recall from the earlier discussion that a number of empirical studies of the innovation production function have suggested that it may indeed be subject to initial increasing returns). So suppose there is some $\bar{x}$ in $(0, \infty)$ such that $\rho(x) = 1$, $x \geq \bar{x}$ and $\rho(x) < 1$, $x < \bar{x}$ and convex.[11] Define the **bold play path** as

$$\begin{aligned} x_t &= \bar{x}, && \text{for } y_t \geq \bar{x} \\ x_m &= y_m, && \text{where } m = \inf\{t:\ y - t\bar{x} < \bar{x}\} \\ x_t &= 0, && t > m. \end{aligned}$$

In words, allocations of $\bar{x}$ (ensuring success) are used when the remaining budget is greater than $\bar{x}$. The first time the budget dips below $\bar{x}$, all of the remainder is exhausted.

Proposition 8 **Suppose that in addition to (A1)′ and (A2), ρ is convex and**

increasing over some initial interval. Then, the bold play path is optimal.

Proof: In Appendix B. □

Remark: A similar result is true for S–shaped ρ, i.e., that for which there are initial increasing returns to allocation, followed by subsequent decreasing returns. In this case, allocations are decreasing over time but greater than $\bar{x}$ as long as the remaining budget is at least as large as $\bar{x}$. Thereafter, everything is staked in one period.

6. Terminal Payoffs

In a pure research project it may be reasonable to suppose that returns accrue only after all stages of the project have been successfully completed. I turn now to this terminal payoff specification. As in the flow payoff model, the analysis is contingent on whether or not the R&D manager can make one or an infinite number of attempts in order to complete any one stage; however, the results are very similar in the two cases. I start with the single attempt specification. Letting W denote the terminal payoffs, the problem of allocating a budget of size y is: $\max W. \prod_{\zeta=1}^{N} \rho(x_\zeta)$ s.t. $\sum_{\zeta=1}^{N} x_\zeta \leq y$.

Proposition 9 **Suppose (A1) holds. Let $(x_1^*, ..., x_N^*)$ be an optimal allocation for the terminal payoff problem. Then,**

i) any permutation $(x_{\sigma_1}^*, ..., x_{\sigma_N}^*)$ of x^* is also an optimal solution. Further, if $\rho(0) = 0$, then $x_t^* > 0$, $t = 1, ..., N$.

ii) if $\ln \rho$ is a concave function, $x_t^* = y/N$, $t = 1, ..., N$, i.e., spreading the budget evenly, is an optimal solution. If $\ln \rho$ is strictly concave, equal allocation in the only optimal allocation.

iii) if $\ln \rho$ is convex over [0,y], then $x_\zeta^* = y$ for some ζ, $x_t^* = 0$, for all $t \neq \zeta$, is an optimal allocation. If $\ln\rho$ is strictly convex, that is the only optimal allocation.

Proof: The proofs are all trivial. For ii) and iii), note only that an equivalent maximizization problem is to maximize the log of the objective function, i.e., max $\sum_{t=1}^{N} \ln\rho(x_t)$ s.t. $\sum_{t=1}^{T} x_t \leq y$, $x_t \geq 0$. □

Suppose instead that any number of attempts can be made till (eventually) a stage is successfully completed. As in the flow payoff model I shall make the simplifying assumption that the **total** allocation for any stage is all that the decision–maker can choose. It is evident that $V(y,n) = \max_{x \in [0,y]} V(y-x,n+1)\tau(x)$. It should be straightforward to see that the infinite mistake problem reduces to max $\prod_{n=1}^{N} \tau(x_n)$ s.t. $\sum_{n=1}^{N} x_n \leq y$. Consider the following assumption on the distribution of completion times whose implication is that the expected discounted time to success, τ, is a concave function:

(A4) $\lambda F(.;x_1) + (1-\lambda)F(.;x_2)$ first–order stochastically dominates $F(.;\lambda x_1+(1-\lambda)x_2)$ for all x_1, x_2 and $\lambda \in [0,1]$.

Remark: If the distribution of completion times is exponential, i.e. $F(T;x) = 1 - e^{-Tx}$, then (A4) is satisfied.

Proposition 10 **Under (A4), the optimal allocation is to distribute the budget evenly between all stages of the project, i.e $x_n^* = y/n$.**

7. Extensions and Discussion

In this section I discuss two extensions and some avenues of future research. The first generalization is to allow probabilities of success that depend not just on the current allocation but also on the current stage. This is motivated by the observation that the stages in the R&D project may refer to very different tasks or problems.[12]

In general, the characterization requirements are going to be more stringent. To understand why, consider the flow payoff problem. With ρ independent of n but increasing the earlier the period the higher the allocation; for any two consecutive periods (t, t+1), allocations x_t, x_{t+1} affect probabilities of success in periods t + 1 and beyond in an identical way but x_t also affects the probability of surviving until $t + 1$. The second argument still holds but the first is modified as follows. Since $\rho(x_t,t)$ and $\rho(x_{t+1},t+1)$ enter multiplicatively in determining the probability of success in periods t+1 onwards, **if** the returns from period t+1 were the only consideration, the optimal allocation would maximize this product, i.e. should equate the marginal returns to $\ln \rho(.,t)$. Given the additional desirability of attaining stage t+1, an optimal allocation in the flow payoff model exhibits the property $\frac{d}{dx}\ln \rho(x_t^*,t) \leq \frac{d}{dx}\ln \rho(x_{t+1}^*,t+1)$ (see Dutta (1988) for details).[13] Indeed the same logic also says that in the terminal payoff specification, the optimal allocations have the property that $\frac{d}{dx}\ln \rho(x_t^*,t) = \frac{d}{dx}\ln \rho(x_{t+1}^*,t+1)$. (A number of other results may be found in Dutta (1988)).

A second generalization is to admit scrap–value payments for abandonement of the project. Qualitatively, the results remain unchanged and indeed in the discussion in previous sections I have already indicated some details on this issue.

A few comments on the interpretation of the models studied in this paper are in order. The models apply directly to the R&D decision–making of a monopolist or a competitive firm which believes that its decisions do not alter the R&D expenditures of its competitors. In the latter case the probability of success function and the stage rewards are determined by the rivals' actions as well as technological factors. If the decision–maker is an R&D manager then the flow payoff model is an analysis of a manager who is paid γ_n only when the n–th stage is completed. The single attempt model is an incentive scheme in which, additionally, the manager loses his job in the event of failure. The terminal payoff model, then, is an incentive scheme in which the manager is paid only when the entire project is successfully completed; again the single attempt version incorporates dismissal for unsatisfactory performance.

I turn now to future research issues. I did not address in this paper the question: why is there an aggregate budget constraint for the R&D allocation problem? I believe one reason has to do with the fact that R&D is typically carried out by research groups within the firm or contracted out; the separation between firm management and R&D management creates an agency problem and a budgetary mechanism is one incentive system that deals with this problem. The premise of the argument is that, ceteris paribus, it is better to give R&D managers flexibility in resource allocation; for example, they may acquire information about the level of difficulty or the time to completion of various stages along the way and the greater the flexibility they have the more profitably (and quickly) they can address this information. However complete discretion on expenditure may have bad incentive effects; R&D managers are prone to waste funds or not take adequate precautions if they do not face a budget or cost constraint. These arguments are, of course, conjectures at this point and it remains to construct a more formal framework within which the aggregate budget specification can be endogeneously derived.[14]

Granot and Zuckerman (1991) and Gallini and Kotowitz (1985) concerned themselves with the question: if the R&D manager has a choice over the sequence in which he picks the stages of research to pursue, what would be an optimal way to do so? Note that in their formulations the manager **does not** choose any expenditure levels but rather has available a menu of immediate rewards γ_n and probabilities ρ_n and selects sequentially from that menu. The techniques employed in this paper may be useful in combining the sequencing and expenditure problems.

Grossman and Shapiro (1986) showed that in the problem with an aggregate budget constraint and terminal payoffs the optimal time–path of allocation is to increase the stage expenditures as the project progresses. This contrasts with my result in Proposition 9 that the optimal allocation is stage–independent. The intuition for their result is of course that, in the absence of a budget constraint, the allocation is determined period by period and the marginal benefits are higher when a greater number of stages have already been successfully completed.

As noted in the introduction, a related literature is that on the "bomber problem" in operations research. Two questions have been addressed in this literature: i) what is the expected survival time for the gunner (this is the flow payoff problem under the restriction that $\gamma_n = 1$ for all n and a single attempt at each stage)? And ii) what is the probability of surviving all of the enemy aircraft (this is the terminal payoff problem with a single attempt at each stage). (See Shepp, Simmons and Yao (1990) for i) and Simmons and Yao (1990) for ii)). There are three main points of difference between this paper and the above literature. First, much of that literature places functional form restrictions on ρ; for example, Shepp, Simmons and Yao (1990) examines the case $\rho(x) = 1 - e^{-x}$ (whereas I consider a general formulation for ρ). In that sense, my results on time and budget monotonicity in the flow payoff model can be viewed as a generalization of their results. Second, I examine the bold play issue which is not addressed by this literature. And, third, I analyze multiple attempts and allow stage–dependent payoffs as well. Ross (1983) also discusses the terminal payoff problem and presents a version of Proposition 9.

8. Conclusion

I examined the dynamics of optimal management of a sequential R&D project. I showed that if the total expenditure is fixed a priori, then the optimal allocation in a development project involves greater expenditures in early stages regardless of the specification of flow payoffs. However, if the project is a basic research type whose payoffs come only at the end, then the optimal allocation is to spread the budget evenly if the innovation production function is subject to decreasing returns. Increasing the size of the budget does not in general lead to higher allocations at every stage although it does increase the budget remaining after each stage. Finally, an Inada–type boundary condition is both necessary as well as sufficient for the optimality of "bold play" in the usage of the budget. All of these results hold in the two alternative specifications of a single permissible attempt at each stage or an infinite number of attempts.

Footnotes

1 In Sections 2–4, Grossman and Shapiro studied a sequential R&D formulation to which there is no immediate analog in this paper.

2 A third set of dynamic considerations relate to the following "stopping" or adoption problem: an R&D manager explores a sequence of technological opportunities and has to determine at which point to stop exploring and adopt a currently available technology. This set of adoption issues are at the heart of Roberts and Weitzman (1981)– and I do not investigate them in this paper. Some game–theoretic models also examine dynamic R&D issues (for example, Fudenberg et.al (1983) and Harris and Vickers (1989)), although their discussion of the dynamics is more limited and their focus more directly on the strategic aspects of the problem.

3 In principle, the expenditure should also be contingent on the number of further attempts that can be made at completing the current stage; in the two cases I study, single attempt and infinite attempts, the firm respectively, cannot and need not condition on this variable and hence I ignore it in the above formulation.

4 The distribution will also depend on which of the two cases of admissible number of mistakes, single or infinite, is being analyzed.

5 The sufficient condition for existence is unduly strong. Weaker conditions are not offered here since their statement would involve additional notation and other unnecessary complications and further the existence issue is not the focus of this paper. However note that Dutta (1988) shows, for example, that in the instance that the stage game payoffs are constant, i.e. $\gamma_n = \gamma$, payoffs are well–defined and there is a solution to the infinite stage problem even when $\delta = 1$.

6 If the scrap–values also depend on the size of unspent budget, the proposition goes through provided $\rho(x)\gamma_n(x) + [1-\rho(x)]S_n(y-x)$ is non–decreasing in x, for all y.

7 Such a systematic increase (or decrease) in the allocation turns out to be equivalent to the property that the cross–partial of the value function be positive (or negative) throughout. As we will see in the sequel it may not even be possible to ensure more elementary properties like the concavity of the value function, leave alone modularity properties.

8 The boundedness condition that suffices for the existence of an optimal solution is satisfied if $\delta\beta < 1$.

9 Assumption (A4) in Section 6 is a sufficient condition under which the expected discounted time to success, τ, is a strictly concave function.

10 The fact that ρ is strictly increasing at $x = 0$ rules out the possibility that $\rho(x) = 0$ over some initial interval of allocations. Indeed that is desirable because bold play would otherwise, and trivially, be optimal behavior as a consequence of that

possibility alone.

[11] (A1') should be taken to mean that ρ is differentiable over $[0,\bar{x})$.

[12] When the probability of success functions are contingent on immediate allocations as well as the stage itself, i.e. they may be denoted $\rho(.,n)$, existence of optimal budgetary strategies in the potentially troubling case where the number of stages is infinity, requires joint restrictions on $\rho(.,n)$, the immediate payoffs γ_n and the discount factor δ. For a fuller discussion of the existence issue, see Dutta (1988).

[13] A natural question to ask is: under what conditions on $\rho(.,n)$ do the actual allocations decline with time? Dutta (1988) presents several sufficient conditions for this to hold; by way of an example, it may be noted that if ρ is multiplicative in stage and allocation, i.e. $\rho(x,n) \equiv k(n)\rho(x)$, then actual R&D expenditures do in fact decline over time.

[14] I have taken the size of the budget to be exogeneously given and investigated its optimal usage. It is evident that one can investigate the question of the optimal size of the budget y by combining the two functions; lifetime profits from optimal allocation, $V(y,1)$, and a cost function, say $c(y)$.

References

Dasgupta, P. and J. Stiglitz (1980). Uncertainty, Industrial Structure and the Speed of R&D, **Bell Journal of Economics** 11, 1–28.

Dubins, J. and L. Savage, 1965. Some Inequalities in the Theory of Stochastic Processes, Dover, New York.

Dutta, P. K., (1988). Budget Allocation Under Sequential Uncertainty, Columbia University Working Paper #381.

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

Gallini, N. and Y. Kotowitz, (1985). Optimal R&D Processes and Competition, **Economica** 52, 321–334.

Granot, D. and D. Zuckerman, (1991). Optimal Sequencing and Resource Allocation in Research and Development Projects, **Management Science**, 37, 140–156.

Grossman, G and C. Shapiro, (1986). Optimal Dynamic R&D Programs, **Rand Journal of Economics**, 17, 581–596.

Harris, C. and J. Vickers, (1986). Perfect Equilibrium in a Model of Race, **Review of Economic Studies**, 54, 1–29.

Kamien, M. and N. Schwartz, 1982, Market Structure and Innovation, Cambridge University Press, Cambridge, England.

Lucas, R. (1971). Optimal Management of a Research and Development Project, **Management Science**, 17, 679–697.

Roberts, K. and M. Weitzman, (1981). Funding Criteria for Research, Development and Exploration Projects, **Econometrica**, 49, 1261–1288.

Ross, S., 1983, An Introduction to Stochastic Dynamic Programming, Academic Press, New York.

Shepp, L., G. Simmons and Y–C Yao, (1990). On a Problem of Ammunition Rationing, mimeo, AT&T Bell Laboratories.

Simmons, G. and Y–C Yao, (1990), Some Remarks on the Bomber Problem, **Advances in Applied Probability**, 22, 412–432.

Stokey, N., R. Lucas and E. Prescott, 1989, Recursive Methods in Economic Dynamics, Harvard University Press, Cambridge, Mass.

Appendix A

In this appendix I provide details of the computations involved in Example 2. Recall from (4.2) that the optimal choice involves

$$\underset{x \in [0,y]}{\text{Max}} \frac{x}{x+1}\left[1 + \left(\frac{y-x}{y-x+2}\right)^2\right] \tag{A.1}$$

Let the maximand in (A.1) be denoted $\varphi(x;y)$. A straightforward computation reveals that

$$\varphi_x = \left[\frac{1}{(x+1)}\right]f(x,y) \tag{A.2}$$

where

$$f(x,y) = \frac{1}{x+1}\left[1 + \left(\frac{y-x}{y-x+2}\right)^2\right] - 4x\left[\frac{y-x}{(y-x+2)^3}\right]$$

Evidently, from (A.2) it follows that φ_x has the same sign as f(x,y) on $x \geq 0$. Further, a little algebra reveals that on $x \in [0,y]$, $f(x,y) \geq 0$ if and only if the function $\psi(x,y) \geq 0$, where

$$\psi(x,y) = x^3 + (y+6)x^2 - (3y^2+10y+6)x + (y^3+4y^2+6y+9) \tag{A.3}$$

Notice that $\psi(0,y) > 0$ as also $\psi(y,y) > 0$. Further,

$$\psi_x(x,y) = 3x^2 + 2(y+6)x - (3y^2+10y+6) \tag{A.4}$$

It is immediate from (A.4) that $\psi_x(0,y) < 0$ and $\psi_x(y,y) = 2y^2 + 2y - 6$. In other words, for all y ψ is a strictly convex function which is strictly positive but declining at x = 0 and strictly positive also at x = y. Moreover, from the above expressions it is clear that for small y, say $y < \tilde{y}$, $\psi > 0$ throughout, i.e. that $\varphi_x > 0$ on [0,y] and hence that the optimal solution of (A.1) is attained at $x^* = y$. For larger y, it follows from (A.4) that ψ is as depicted in Figure 3:

(Figure 3)

Remembering that the sign of ψ is the same as that of φ_x it is immediate that, for such y, φ increases over $[0,x_1(y)]$, decreases over $[x_1(y),x_2(y)]$ and increases again over $[x_2(y),y]$. Hence, the maximum payoff is achieved either at $x_1(y)$ or at y.

Claim: For large enough y, the optimum choice must be at $x_1(y)$. To see this note that $\varphi(y/2;y) = \frac{y}{y+2}[1 + (\frac{y}{y+4})^2]$, whereas $\varphi(y;y) = \frac{y}{y+1}$. Simple algebraic manipulation reveals that there is a critical y' such that $\varphi(y/2;y) > \varphi(y;y)$ iff $y > y'$. The claim follows.

We know then, that for $y < \tilde{y}$, the optimal choice is $x^* = y$ and that for $y > y'$, the optimal choice is at $x^* = x_1(y) < y$. Furthermore, by examining the conditions defining $x_1(y)$, it is straightforward though tedious to show that there cannot be an interval over which $\varphi(y;y) = \varphi(x_1(y);y)$; since the correspondence of maximizers to (A.1) is upper semi–continuous, there is clearly **one** such point at which the decision–maker is indifferent between $x_1(y)$ and y. Collecting all of this, we can assert that there is some $\hat{y}$ to the left of which the unique optimal choice is $x^* = y$ and to the right of which the unique optimal choice becomes $x_1(y)$. Since, $x_1(y) < y$, the optimal allocation "jumps down" at $\hat{y}$.■

Appendix B

Proof of Proposition 3: I first demonstrate that the value function has the form claimed above. Consider being at n completed stages with a remaining budget of y. A candidate budgetary strategy is to mimic the optimal allocation that is employed from n+1 completed stages and an identical budget. Clearly, the returns to doing so are $1/\beta\, V(y,n+1)$ and this, by definition is no greater than $V(y,n)$, i.e. $V(y,n+1) \leq \beta V(y,n)$. However, this argument inverts; starting from n+1 completed stages with a remaining budget of y, a feasible budgetary strategy is to mimic the optimal strategy from n completed stages and an identical leftover budget. That argument yields, $V(y,n+1) \geq \beta V(y,n)$. The value function evidently has the claimed structure.

The optimality equation for this problem is

$$V(y,n) = \max_{0 \le x \le y} \{\rho(x)\gamma_n + \delta[\rho(x)V(y-x,n+1) + (1-\rho(x))V(y-x,n)]\}$$

Given the form of the value function, writing V(y) for V(y,0) and $\gamma\beta^n$ for γ_n, the optimality equation reduces to

$$V(y,n) = \beta^n \max_{0 \le x \le y} \{\rho(x)\gamma + \delta[\rho(x)\beta V(y-x) + (1-\rho(x))V(y-x)]\}$$

Evidently, the maximization is independent of n and hence so is any optimal budget strategy (which are all selections from the set of maximizers). □

Proof of Proposition 8: Let me first show that if $y \le \bar{x}$, then immediate exhaustion is optimal. Since this is evidently true for stage n=N, suppose in fact that n < N. So I claim that at such a stage the value function for budgets in $[0,\bar{x}]$ is $V(y,n) \equiv \rho(y)[\gamma_n + W_n]$, where $W_n = \sum_{t=1}^{N-n} (\delta\rho(0))^{t-1} \gamma_{n+t}$. Clearly it suffices to check that this function satisfies the optimality equation over $[0,\bar{x}]$, i.e.

$$V(y,n) = \max_{x \in [0,y]} \left[\rho(x) + \delta\rho(x)\rho(y-x)W_{n+1}\right], \forall\, y \le \bar{x} \tag{A.5}$$

It is easy to show that $\rho(x)\rho(y-x)$ is a convex function of x. Hence the maximand in (A.5) is convex and so the maximum is achieved at either 0 or y. From Proposition 1 it is known that the maximum cannot be achieved at 0 (else, $x_t = 0$ for all t thereafter and that allocation cannot be optimal).

Suppose now that we start with $y > \bar{x}$. By the above arguments the budget is always exhausted after a finite number of allocations even if there are infinite stages in the R&D project. The last positive allocation is in $(0, \bar{x}]$. Without loss of generality we can restrict attention to the infinite stage problem; for the finite stage problem if the last positive allocation is at least $\bar{x}$ then so are the previous ones and by Proposition 1 the bold play path is optimal while on the other hand if the last allocation is in $(0, \bar{x}]$, the analysis is identical to that which now follows. I shall now show that the penultimate positive allocation is $\bar{x}$. Clearly I only need to show that an allocation as $(x_1, x_2, 0, 0...)$ where $\bar{x} > x_1 \geq x_2 > 0$ and $x_1 + x_2 > \bar{x}$, is inoptimal. The argument is identical to that used for the case above. Note that the maximization problem is the same as (A.5) except for the fact that the effective domain from which x is chosen is $[0,\bar{x}]$. Since the maximand is convex the maximizer must be either of the two extremes and it cannot be 0. The proposition is proved. □

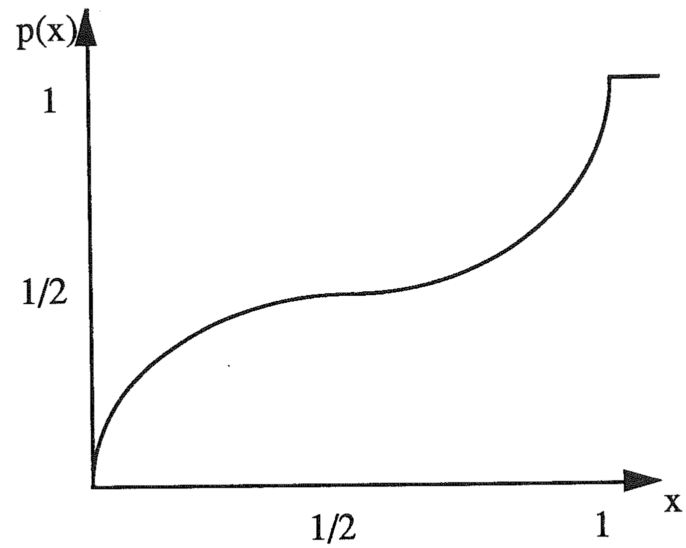

FIGURE 1

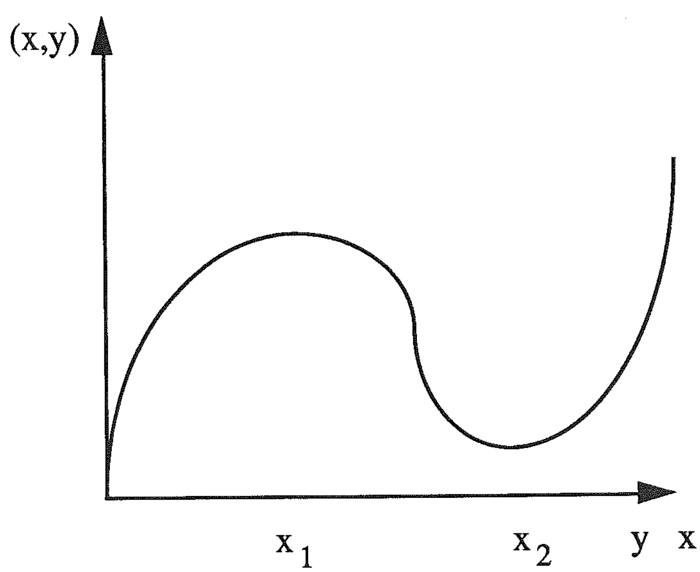

FIGURE 2

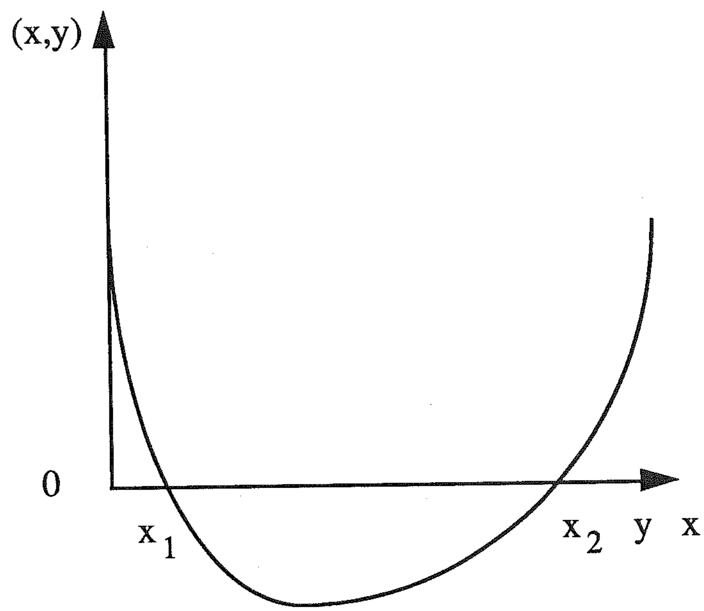

FIGURE 3