Search and the Introduction of Improved Technologies

Laura J. Kornish,[1] Steven A. Lippman,[2] John W. Mamer[2]

[1] *Leeds School of Business, University of Colorado, Boulder, Colorado 80309-0419*

[2] *The John E. Anderson School of Management at UCLA, Los Angeles, California 90095-1481*

Received 31 August 2010; revised 27 April 2011; accepted 28 April 2011
DOI 10.1002/nav.20468
Published online 30 June 2011 in Wiley Online Library (wileyonlinelibrary.com).

Abstract: Modeling R&D as standard sequential search, we consider a monopolist who can implement a sequence of technological discoveries during the technology search process: he earns revenue on his installed technology while he engages in R&D to find improved technology. What is not standard is that he has a finite number of opportunities to introduce improved technology. We show that his optimal policy is characterized by thresholds $\xi_i(x)$: introduce the newly found technology if and only if it exceeds $\xi_i(x)$ when x is the state of the currently installed technology and i is the number of remaining introductions allowed. We also analyze a nonstationary learning-by-doing model in which the monopolist's experience in implementing new technologies imparts increased capability in generating new technologies. Because this nonstationary model is not in the class of monotone stopping problems, a number of surprising results hold and several seemingly obvious properties of the stationary model no longer hold.
 Naval Research Logistics 58: 578–594, 2011

Keywords: search; technological improvement; learning by doing

1. INTRODUCTION

The decision to release a new version of a product involves a delicate balance between often conflicting forces. Whereas, the increased quality (and therefore market position) bestowed by a new product introduction improves the firm's profit stream, releasing a new product reduces the number of future product improvements that can be made. The constraint of "product architecture" is most apparent in technology products. Each improvement to a software product that maintains compatibility with extant users increases the cost and complexity of supporting and maintaining the product. Eventually, the product, though vastly improved from multiple improving releases, becomes unwieldy and must be rewritten and replaced. Similar considerations apply to digital hardware: adding features eventually forces the device to be "re-engineered." Combined with customer resistance to a sequence of product introductions, architectural limitations render a new product introduction opportunity a scarce resource to be carefully managed.

This article focuses on two trade-offs. The first entails the balance between immediately reaping the reward from an improvement and deferring the reward in favor of an even bigger improvement (and concomitant reward) in the future. The second trade-off balances the ongoing cost of R&D with the anticipated rewards from R&D. When the prospects for improvement are dim, it makes sense to stop searching for the "new new thing"[1].

The structure of this problem and the associated trade-offs apply to any innovation-intensive endeavor in which a sequence of improved versions can be introduced. In addition to improving software, the "introduce or wait" dilemma applies to increasingly fast computers, increasingly compact cellular phones, and increasingly sophisticated network products. Process improvements are subject to the same logic. An implement-each-improvement in factory configuration or work environment approach fails to anticipate future improvements that may surface: it fails to balance the first trade-off.

We frame the decision problem as a multi-stage dynamic optimization problem with sequential introductions under

Correspondence to: J.W. Mamer (jmamer@anderson.ucla.edu)

[1] Operating in a competitive environment introduces a third trade-off: get to market early, perhaps with a minor innovation, to secure a dominant position vs. waiting to secure a more substantial innovation at the risk of being pre-empted by competitors. This trade-off has been studied in a winner-take-all context by Lippman and Mamer [18] and is not considered in this article.

uncertainty in which a firm can introduce up to N (additional) innovations. Finite N captures the essential constraint on the number of improving introductions that can be made—brought about by either the product architecture, the willingness of customers to accept new versions, or even an endogenously specified product life cycle.[2] We model the discovery of the sequence of technological innovations as a compound Poisson process: the times between arrivals are exponentially distributed random variables, and the size of the innovations are independent random draws from a single (unchanging) distribution. This compound Poisson process captures uncertainty in both the timing and the size of the innovations.

Our analysis addresses the two trade-offs: the conflict between something good now and something better later (now vs. later) and the balance between the potential for improvement and the cost of finding it (continue vs. stop). We show in Theorem 1 that, like other search models, the optimal policy is characterized by thresholds: adopt the new technology if and only if its value exceeds $\xi_i(x)$, where x is the state of the technology currently in place and i is the number of remaining implementations allowed. The threshold $\xi_i(x)$ is the solution to the now vs. later trade-off: the firm should wait to introduce a new product until a sufficiently large improvement has been found. Theorems 2 and 3 demonstrate that these thresholds are increasing in the state x of the most recently implemented innovation and decreasing in the number i of possible introductions remaining.

Theorem 1 also shows that a single threshold solves the continue vs. stop trade-off: there is an innovation level $\overline{\xi}$ such that it is optimal for the firm to engage in R&D if and only if the size x of the most recently implemented technology is less than $\overline{\xi}$. If the firm discovers an innovation above this threshold at any point in time, no matter how many possible introductions (out of the N) are remaining, it is introduced and search ends.

Based on the idea of learning-by-doing (a firm's innovative capability increases with the number of innovations it has introduced), we extend the model by examining a learning effect that takes place when the firm introduces an innovation. The traditional learning curve is a relationship between marginal cost of production and cumulative production. In our context, we relate the parameters that characterize the search environment to experience with introductions. We examine the effects on the optimal policy when the introduction of a new technology improves future search. Like learning-by-doing, experience improves R&D capability because the firm has a chance to learn from mistakes and problems encountered in the introduction of new technological discoveries. Theorems 5 and 6 show how an increase in the ability to generate technological discoveries impacts the optimal policy. As generating ability increases, through some combination of making innovations come faster, reducing the cost of search, or improving the distribution of innovations, the threshold $\xi_i(x)$ decreases in x for $i > 1$ if the expected duration of search in the last period also decreases. However, if the increased generating ability increases the expected duration of search (e.g., via a reduction in the cost of search per unit time), then the current threshold can increase or decrease.

While our research is directed at the problem of technology adoption, our search model has other applications. For example, consider a professional worker who engages in "on-the-job" search: while working, he continues to receive job offers. If he switches jobs too often, he risks being labeled a "job hopper", tarnishing his reputation and reducing the quality of future job offers. In this example, it is reasonable to model the worker as limiting himself to a small number N of job switches during his career.

A discussion of related literature and a formal presentation of the model are given in Sections 2 and 3. We analyze the stationary model in Section 4. Section 5 analyzes the non-stationary model induced by learning. A few conclusions are proffered in Section 6.

2. RELATED LITERATURE

The idea that technological change not only leads to economic growth but also is the most important factor accounting for economic growth is now a commonplace. According to Mansfield [23; p. 4, 5, 7], the vast majority of the long-term increase in output per capita in the United States is attributable to technological change along with increased educational levels:

> "advance of knowledge" contributed about 40 percent of the total increase in national income per person employed during 1929–1957. … In most industries, new products account for a significant share of the market.

The accepted notion that repeated innovation is a major driving force in modern economies has roots in the writings of Schumpeter [32] and Usher [34]. Schumpeter describes the importance of the entrepreneur-led "gales of creative destruction" for economic progress in a capitalist society. He distinguished inventions (the flashes of insight) from innovations (the improved ability to profit from the invention) whereas Usher's labels for these two modes of technological improvement are "critical insight" and "critical revision."

[2] Some of these features, for example increasing customer resistance to successive product releases, could have been modeled with a fixed cost of implementing that rises with each new release. This approach would introduce a non-stationary element into the model with a concomitant increase in analytical complexity. Allowing a finite number of new product releases captures the spirit of the product improvement problem while at the same time allowing us to characterize optimal policies.

Gomory [16] contrasts the high-profile process of turning breakthrough scientific discoveries into commercial products (Schumpeter's invention) with the less dramatic and grueling process of incremental product improvement (Schumpeter's innovation) and stresses the importance of the process of refinement. But perhaps it is hindsight, and lack of historical knowledge, that induces us to view the difference between invention and innovation as so brightly illuminated. For example, Scherer [31] discusses this process of revision in the context of James Watt's activities from 1765–1780 improving the steam engine; he writes (pp. 21–22), "Many [improvements] were logical extensions of Watt's original idea, requiring relatively little new insight but considerable trial and error. In other instances … a high level of creative insight was revealed." Providing a multitude of examples, economic historian Rosenberg [29, p. 7] notes that

> inventive activity is, itself, best described as a gradual process of accretion, a cumulation of minor improvements, modifications, and economies, a sequence of events, where, in general, continuities are much more important than discontinuities. …even the big technological breakthroughs which are associated with the names of Darby, Watt, Cort, and Bessemer, usually have much more gently declining slopes of cost reduction flowing from their technical contributions than the historical literature would lead us to expect.

One consequence of the "accretion …of minor improvements" is that all R&D effort does not lead to improvements; some attempted improvements will not be fruitful. In our model, R&D effort is synonymous with search; because these efforts consist of independent draws from an unchanging distribution (in which each technological improvement is completely characterized by its level of profitability), the spirit of our article is aligned with the incremental improvements of Usherian revision. In this context, we contribute to the literatures of technology choice and search theory by exploring the case of repeated search for and adoption of improved alternatives.

In the standard search problem, the searcher can continue taking independent draws from F, the unchanging distribution of the size of newly found technological innovations; the size of an innovation is simply the per unit time revenue associated with the innovation. On finding an acceptable innovation, the monopolist implements the innovation and ceases search. Lippman and McCall [19] describe many variations on this standard search problem. The closest variation to the problem we study here is the on-the-job search model (Lippman and McCall [19], pp. 179–181; Burdett [5]). In on-the-job search, the decision maker chooses between working only (and not searching), searching only (and not working), and searching while working. Employed searchers accept any offer that increases their wage and quit searching at very high wage levels. This latter aspect of the solution is similar to the resolution of the continue vs. stop trade-off in our model. However, in contrast to our work, the focus in the on-the-job search models is whether or not to search while working and not which job offers (innovation sizes in our model) to accept.

Articles that explore this question of which innovations to implement when facing a stream of improved innovations include Balcer and Lippman [3], Cauley and Lippman [7], and Farzin et al. [15].[3] In these three articles as well as this article, the optimal policy is characterized by one or more thresholds: implement when the improved innovation exceeds the threshold. The work of Balcer and Lippman [3] is formulated as a repeated purchase decision problem, whereas our model as well as the models in these other two articles are couched in terms of a firm's repeated decision of whether or not to introduce a new technological discovery. Like this work, the models in all three articles are set in an environment of uncertainty, with uncertain time until innovation and uncertain innovation size. In our model and the model of Farzin et al., introduction of new innovations is discouraged due to a limit on the number of introductions. The other two articles do not impose a limit on the number of introductions; instead, introduction of new innovations is discouraged by switching costs.

The work of Cauley and Lippman [7] models endogenously generated technological discoveries in which the level of a firm's currently installed technology affects the distribution of technological discoveries; moreover, the efficiency of the R&D effort increases with the state of technology (as could be generated with learning-by-doing). In the articles by Balcer and Lippman and Farzin et al., the efficiency of the R&D effort is unchanging both in time and in the state of technology. In this article, however, the efficiency of the R&D effort declines with time because the parameters of the search process are stationary: search for technological innovations is drawn from an unchanging offer distribution. The case of non-stationary parameters is treated in the penultimate section.

In these three articles, the search process produces a sequence of strict improvements. This contrasts with our work: the value of the technological discoveries, drawn from the same distribution, can fall below the value of the technology currently in place. Moreover, because the distribution of innovations is time invariant, the economic return to search decreases with the state of the technology currently in place whence eventually the monopolist ceases search. The three works cited above that model guaranteed improvements as

[3] Also see the follow-up notes on Balcer and Lippman by Kornish [17] and on Farzin et al. by Doraszelski [14]. Repeated technological improvements are not unlike repeated tree-cutting models—for example, see Miller and Voltaire [25, 26]—in which an improving asset can be periodically harvested.

a result of R&D effort fit Schumpeter's invention setting whereas technological improvement in our article is more akin to Schumpeter's innovation.

Gomory's advocacy of the importance of product refinement rests on a learning-by-doing premise. This premise was advanced by Arrow [2] in the context of manufacturing efficiencies that improve with cumulative production. More recently, Cohen and Levinthal [11, 12] coined the term "absorptive capacity" to refer to a firm's ability to assimilate and exploit knowledge created outside the firm. Investment in absorptive capacity allows the firm to attain a deeper understanding of new knowledge created in the pertinent domain and also to better profit from it. In an extension of the basic model, we explicitly capture the notion of absorptive capacity by modeling an increase in knowledge that emanates from the monopolist's implementation of technological discoveries: implementation itself generates an increase in the monopolist's future ability to uncover profit improving innovations. In Section 5, we explore how this aspect of learning-via-implementation impacts the firm's decisions.

3. MODEL

As is common in the literature, we model the R&D process as sequential search. In this article, we are concerned with a monopolist's decisions regarding which of his newly found innovations, equivalently discoveries, to implement. Although there is no limit on how often the monopolist can sample the distribution of possible innovations, he is limited to a finite number N of implementations. Because $N < \infty$, the monopolist will not elect to implement each and every improved technology he finds. Customer resistance and engineering constraints are two important impediments to frequent product changes. Additional motivation for a fixed N of moderate size was given in the introduction: architectural constraints or increasing complexity can hamper implementing more than a few product improvements, or a limited number of introductions can be imposed by management fiat in order to protect the product identity, or to impose a manageable discipline on the product introduction process[4].

While engaged in R&D, the monopolist expends an amount of money $c \geq 0$ per unit time on R&D: the cost of search is c per unit time. When the monopolist stops expending money on R&D, the discovery rate falls to 0. While pursuing R&D, technological discoveries arrive in accord with a Poisson process with arrival rate λ, and the net revenue flow per unit time of the newly discovered technologies (if implemented) are independent with cumulative distribution function F. To avoid unnecessary complexity, we assume that F has a density f and that the set where f is positive is a (possibly infinite) interval (a, b). All costs and revenues are discounted at the continuous time rate $\alpha > 0$. The time horizon is infinite, and the monopolist's goal is to maximize his expected discounted stream of revenue net of search costs.

If the revenue generated by the technology currently in place is x per unit time, we say that x is the state of the system. If no technology has been implemented, then the state of the system is 0. The monopolist's cost per unit time is 0 if he is not engaged in R&D, and it is c if he is engaged in R&D. When $N - i$ technologies have already been implemented, so that the monopolist is allowed to implement at most i additional technologies, we say that i periods remain or we are in period i. If 0 periods remain, then no additional technologies can be implemented (and the monopolist necessarily ceases search). Consequently, when 0 periods remain and the state of the system is x, the monopolist's future expected discounted profit is simply x/α. The questions to be answered are when to cease search and which technologies are worthy of introduction.

Although it might be natural to require that any adopted technology be a strict improvement over the one in place, we do not impose this requirement; nevertheless, it falls out of our analysis. Similarly, we allow recall of any previously discovered but not implemented technology to simplify the presentation; again, recall of past discoveries is never exercised.

4. ANALYSIS

In the standard continuous time infinite horizon economic search problem with no discounting and Poisson arrival of offers at rate λ, the searcher expends c per unit time until he encounters an acceptable offer, say of value x. On accepting this offer, he receives x and the problem terminates. The optimal policy is characterized by a threshold ξ called the reservation price: accept the first offer whose value is at least ξ. The threshold ξ is the unique solution to (see Lippman and McCall, [19])

$$c = \lambda H(y) \tag{1}$$

where

$$H(y) = \int_{y}^{\infty} (s - y) f(s) ds. \tag{2}$$

Moreover, ξ is the value of the problem. In the classical model, the optimal policy is "myopic": the searcher stops

[4] An alternative would be to propose a fixed cost for introducing a new product. Necessarily, this fixed cost would have to change with time to capture, for example the increasing complexity of the modified product. We imagine the results of such a model would be similar to those of this article; however, we believe that modeling the constrained plasticity of a product as a finite number of introduction opportunities best captures the product improvement situation.

(optimally) when the incremental discounted expected value of exactly one more observation falls below the discounted expected cost of obtaining that observation.

Now presume that all costs and revenues are discounted at rate $\alpha > 0$. It is a trivial extension, but useful for our purposes, to let an offer of x mean that if it were accepted, then x per unit time would be received from that moment onward. Equivalently, on accepting an offer of x, the searcher receives a lump sum of x/α. If he were to continue search until the next offer arrives, with the ability to recall the extant offer of x, then his expected return would be

$$-\frac{c}{\alpha+\lambda}+\frac{\lambda}{\alpha+\lambda}\left[xF(x)+\int_x^{\infty} sf(s)ds\right]\Big/\alpha$$
$$=-\frac{c}{\alpha+\lambda}+\frac{\lambda}{\alpha+\lambda}[x+H(x)]/\alpha.$$

When y is the optimal threshold, the searcher is indifferent between accepting y, earning y/α, and taking exactly one more observation. Hence,

$$\frac{y}{\alpha}=\frac{-c}{\alpha+\lambda}+\frac{\lambda}{\alpha+\lambda}[y+H(y)]/\alpha \tag{3}$$

which reduces to

$$y+c=\lambda H(y)/\alpha. \tag{4}$$

The value of y which solves (4) is unique, and we denote it by ξ. Clearly, ξ/α is the value of the problem.

In the standard search problem with discounting, the monopolist is earning revenue 0 per unit time prior to his first and only implementation of some technology. Suppose, however, that he were earning x per unit time when $N = 1$—prior to his first (and last) implementation of some yet to be discovered technology.

In anticipation of our full model, let $V_1(x, y)$ denote the expected discounted return when at most $N = 1$ new innovations can be introduced, the current flow rate of earnings is x, and search continues until an innovation with flow rate at least y is discovered and implemented. Using $\int_y^{\infty} xf(x)dx = H(y) + y\overline{F}(y)$, we have

$$V_1(x,y)=\frac{\left\{x-c+\lambda\int_y^{\infty}\frac{s}{\alpha}f(s)ds\right\}}{(\alpha+\lambda\overline{F}(y))}$$
$$=\frac{y}{\alpha}+\frac{\lambda H(y)/\alpha-(c-x+y)}{(\alpha+\lambda\overline{F}(y))}. \tag{5}$$

The first-order condition is

$$y+c-x=\lambda H(y)/\alpha. \tag{6}$$

Of course, (6) is the same as (4) when $x = 0$. For $x \le \overline{\xi}$, where $\overline{\xi}$ is the unique solution to $c = \lambda H(y)/\alpha$, let $\xi_1(x)$

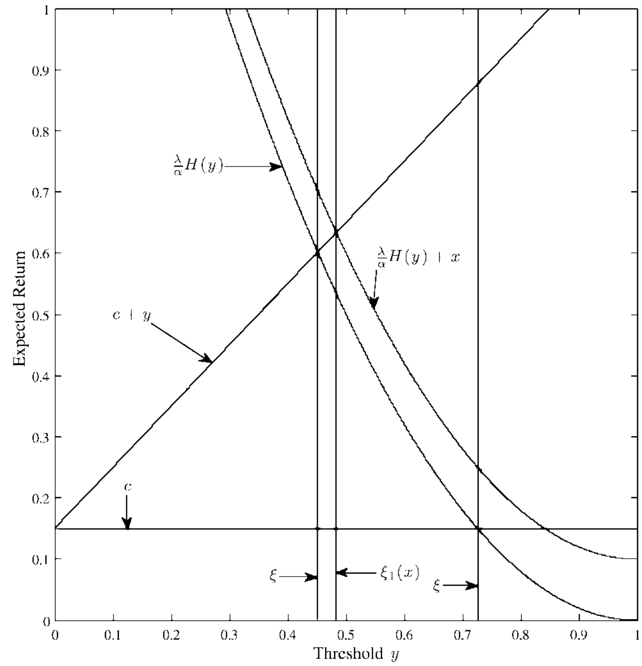

Figure 1. The Calculation of $\xi_1(0.1)$ when $\lambda = 1$, $\alpha = 0.25$, $c = 0.15$, F is uniform on $(0, 1)$.

denote the unique solution to (6). Substituting $y = \xi_1(x)$ into (5) and using (6) produces

$$V_1(x)\equiv V_1(x,\xi_1(x))=\xi_1(x)/\alpha.$$

As made evident in Fig. 1, $\xi_1(x)$ is strictly increasing in x for $x < \overline{\xi}$: as x increases the searcher becomes less willing to implement any new technology. Also, $\xi_1(\overline{\xi}) = \overline{\xi}$ as $V_1(x, y) < x/\lambda$ for all y if $x > \overline{\xi}$: do not search when $N = 1$ and $x \ge \overline{\xi}$. This is the essence of what we find in the model examined in this article: the threshold increases as the current state x increases, and search ceases when the state reaches $\overline{\xi}$.

We will demonstrate that it suffices to limit consideration to threshold policies: every optimal policy is characterized by thresholds. In particular, we let $\xi_i(x)$ denote the optimal threshold when x is the currently installed technology and i implementations remain. See Fig. 2 for a numerical illustration.

Before proceeding with this demonstration, we note four standard facts which we use repeatedly. First, if τ is an exponential random variable with parameter λ, then $Ee^{-\alpha\tau} = \lambda/(\alpha + \lambda)$. Second, if a flow rate c is received at each instant of time for a period whose duration is an exponential amount of time τ with parameter λ, the expected discounted value of these payments is $c/(\alpha + \lambda)$. Third, if each arrival of a Poisson process with parameter λ is recorded (independently of the Poisson process and independently of each other) with probability $\overline{F}(y)$, then the process of recorded arrivals is a

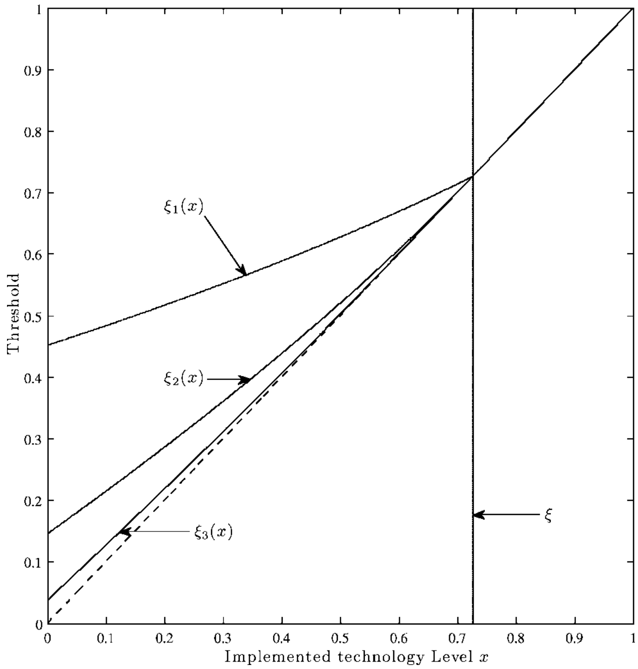

Figure 2. Graph of $\xi_1(x), \xi_2(x)$, and $\xi_3(x)$ when $\lambda = 1$, $\alpha = 0.25$, $c = 0.15$, F is uniform on (0, 1).

Poisson process with parameter $\lambda\overline{F}(y)$ (Ross 1970, p. 29). Fourth, monotone stopping: Let S be the set of states such that it is better to stop than to take exactly one more observation. If S is closed (i.e., once a state in S is reached, all future states of the system also will be in S), then it is optimal to stop at state x if and only if x is in S. Thus, if it is does not pay to take exactly one more observation, then it does not pay to continue sampling for an amount of time τ, where τ is any stopping time for the strong Markov process.[5] The power of this result is that the complex decision of whether or not to stop is reduced to a myopic calculation, stop versus take exactly one more obervation. This result applies to a broad class of stopping problems. It is due to Derman and Sacks [13] and to Chow and Robbins [8] and also can be found in Chow et al. [9]. Also see Lippman and McCall [20, pp. 223–225] for a less abstract treatment.

To ensure that engaging in R&D is a profitable undertaking at the beginning, if not later on, we assume throughout this article that

$$\lambda E(X)/\alpha \geq c. \tag{7}$$

[5] Some boundedness conditions, such as uniform integrability, are required. Application of the monotone stoppping theorem requires verification that S is closed. This is easy to do in most search problems when recall of previous offers is allowed.

4.1. Threshold Policies and Monotonicity of the Optimal Return Function

Our first result states that the optimal policy is completely characterized by thresholds and that there is a critical number $\overline{\xi}$ such that it is optimal for the monopolist to engage in R&D if and only if the state x of the currently installed technology is less than $\overline{\xi}$. Furthermore, the optimal threholds are bounded below by the current state and above by the critical number $\overline{\xi}$.

Let $V_i(x)$ denote the expected return under an optimal policy when i implementations remain and a technology of value x has just been implemented, and define $V_i(x, y)$ to be the expected return to search when i implementations remain, x is the state of the currently installed technology, the monopolist uses threshold y rather than $\xi_i(x)$ in period i, and he acts optimally when fewer than i implementations remain. Clearly,

$$V_i(x) \equiv V_i(x, \xi_i(x)). \tag{8}$$

LEMMA 1: When $i \geq 0$ implementations remain and the current state is x, the optimal return $V_i(x)$ is strictly increasing in x.

PROOF: When 0 implementations remain, $V_0(x) = x/\alpha$. Suppose that $i \geq 1$. Let τ be the time until technology x is replaced when following an optimal policy and the state is x when i periods remain, and let Y_x be the magnitude of the first replacement so that $V_i(x) = E[(x - c)(1 - e^{-\alpha\tau})/\alpha] + E[e^{-\alpha\tau}V_{i-1}(Y_x)]$. Let $\tilde{V}_i(x')$ be the return when starting in state x' when i periods remain, the first replacement occurs at time τ, the new technology has magnitude Y_x, and an optimal policy is followed when $i - 1$ periods remain. Fix $x' > x$ so that $V_i(x') \geq \tilde{V}_i(x') = E[(x' - c)(1 - e^{-\alpha\tau})/\alpha] + E[e^{-\alpha\tau}V_{i-1}(Y_x)] = E[(x' - x)(1 - e^{-\alpha\tau})/\alpha] + V_i(x) > V_i(x)$. Thus, $V_i(\cdot)$ is strictly increasing. □

THEOREM 1: When $i \geq 1$ implementations remain and the current state is x, there is a threshold $\xi_i(x)$ such that an innovation with value y should be implemented if and only if $y \geq \xi_i(x)$. Moreover, $x \leq \xi_i(x) < \overline{\xi}$ for $x < \overline{\xi}$, where $\overline{\xi}$ satisfies[6]

$$c = \frac{\lambda}{\alpha}H(x); \tag{9}$$

$\overline{\xi}$ is the technology level which extinguishes search: for each $i \geq 1$, it is optimal to engage in R&D if and only if $x < \overline{\xi}$.

[6] If $c > 0$, then (9) has a unique solution and $F(\overline{\xi}) < 1$. If $c = 0$, then the existence of a solution to $c = \frac{\lambda}{\alpha}H(x)$ requires that $F(t) = 1$ for some $t < \infty$. In this case set $\overline{\xi} = sup\{t : F(t) < 1\}$. In either case the density f is strictly positive on an interval of the form $(a, \overline{\xi})$. This latter fact is used in the proof of Theorem 2.

PROOF: To show that the optimal policy is characterized by thresholds, suppose i implementations remain, the state is x, the monopolist has discovered (but not yet implemented) a technology y, and $y' > y$. If it is optimal to implement y, then $V_{i-1}(y') \geq V_{i-1}(y) \geq V_i(x)$ where the first inequality follows from monotonicity and the second from it being optimal to implement y. Hence, it is optimal to implement y'. Thus, thresholds characterize the optimal policy.

To demonstrate that $\xi_i(x) \geq x$, we will argue that implementing a technology with value less than the one currently in place can not be optimal. Suppose it were not true, then there is an i and an x such that $\xi_i(x) < x$. Suppose further that the value y of the next technology to arrive when i periods remain and the state is x satisfies $\xi_i(x) < y < x$. By assumption, the monopolist implements y. Consider the policy in which the monopolist does not implement y but thereafter engages in search and implements a new technological innovation if and only if the optimal policy does. Then the net return of this new policy is precisely that of the optimal policy except that it exceeds that of the optimal policy by $x - y$ per unit time during the interval of time starting when the optimal policy implements y and ending when the optimal policy implements a new technology (an interval of exponential length with parameter $\lambda\overline{F}(\xi_{i-1}(y))$). In view of this contradiction, it follows that $\xi_i(x) \geq x$.

Having established that the optimal policy is characterized by thresholds $\xi_i(x)$ and that $\xi_i(x) \geq x$, we now establish that it is optimal to search only when $x < \overline{\xi}$. Observe that

$$V_1(x,y) = \left[x - c + \lambda\int_y^\infty (s/\alpha)f(s)ds\right]\bigg/ (\alpha + \lambda\overline{F}(y)). \tag{10}$$

Differentiating, $V_1(x, y)$ with repect to y, we have

$$\begin{aligned}\frac{\partial V_1(x,y)}{\partial y} &= \frac{\lambda f(y)}{(\alpha+\lambda\overline{F}(y))^2}\left[-(x/\alpha)(\alpha+\lambda\overline{F}(y)) + x - c + \lambda\int_y^\infty (s/\alpha)f(s)ds\right]\\ &= \frac{\lambda f(y)}{(\alpha+\lambda\overline{F}(y))^2}[x - y + (\lambda/\alpha)H(y) - c].\end{aligned}$$

Fix $y > x \geq \overline{\xi}$. Because $H(\cdot)$ is a strictly decreasing function on the support of F, $(\lambda/\alpha)H(y) - c < (\lambda/\alpha)H(\overline{\xi}) - c = 0$ so that $\partial V_1(x,y)/\partial x < 0$ for $y > x \geq \overline{\xi}$ whence, $V_1(x,y) < V_1(x,x)$. Similarly,

$$\begin{aligned}V_1(x,y) &< V_1(x,x)\\ &= \left[x - c + \lambda\int_x^\infty (u/\alpha)f(u)du\right]\bigg/ (\alpha+\lambda\overline{F}(x))\\ &= [x - c + (\lambda/\alpha)x\overline{F}(x) + (\lambda/\alpha)H(x)]/(\alpha+\lambda\overline{F}(x))\\ &= x/\alpha + [-c + (\lambda/\alpha)H(x)]/(\alpha+\lambda\overline{F}(x))\\ &\leq x/\alpha.\end{aligned}$$

Thus, terminating search is preferred to searching when one implementation opportunity remains and $x \geq \overline{\xi}$. Thus, we have established that $V_1(x) = x/\alpha$ for $x \geq \overline{\xi}$.

Assume $V_k(x) = x/\alpha$ for $k = 1, \ldots, n-1$ and $x \geq \overline{\xi}$. Then for $y > x \geq \overline{\xi}$, we have

$$\begin{aligned}V_n(x,y) &= \left[x - c + \lambda\int_y^\infty (s/\alpha)f(s)ds\right]\bigg/ (\alpha+\lambda\overline{F}(y))\\ &= V_1(x,y) < x/\alpha.\end{aligned}$$

Hence, terminating search is preferred when $x \geq \overline{\xi}$.

Finally, we verify that $\xi_i(x) < \overline{\xi}$ for $x < \overline{\xi}$. Fix $x < \overline{\xi}$ and $i \geq 1$, suppose $\xi_i(x) \geq \overline{\xi}$, and assume that the next arrival is at time t and has value y with $\overline{\xi} \leq y \leq \xi_i(x)$. Discounting revenues and costs back to time t, we claim that the return y/α to implementing y at time t exceeds the return of the supposedly optimal policy. Using $\int_w^\infty sf(s)ds = H(w) + w\overline{F}(w)$, we see that the return of the optimal policy is given by

$$\begin{aligned}&\frac{x-c}{\alpha+\lambda\overline{F}(\xi_i(x))} + \frac{\lambda\overline{F}(\xi_i(x))}{\alpha+\lambda\overline{F}(\xi_i(x))}\left[\frac{H(\xi_i(x))}{\alpha\overline{F}(\xi_i(x))} + \frac{\xi_i(x)}{\alpha}\right]\\ &\qquad = \frac{x - c + \xi_i(x)\lambda\overline{F}(\xi_i(x))/\alpha + \frac{\lambda}{\alpha}H(\xi_i(x))}{\alpha+\lambda\overline{F}(\xi_i(x))}.\end{aligned}$$

Define Δ to be the difference between the returns of accepting y immediately and the supposedly optimal policy. Using (9) to substitute for c, we obtain (of course, $y - x > 0$ and $y \geq \overline{\xi}$)

$$\begin{aligned}&\Delta\frac{\alpha}{\lambda}(\alpha+\lambda\overline{F}(\xi_i(x)))\\ &= \frac{\alpha}{\lambda}(y-x) + \overline{F}(\xi_i(x))[y-\xi_i(x)] + [H(\overline{\xi}) - H(\xi_i(x))]\\ &> [\overline{\xi} - \xi_i(x)]\overline{F}(\xi_i(x)) + [H(\overline{\xi}) - H(\xi_i(x))]\\ &= [H(\overline{\xi}) + \overline{\xi}\overline{F}(\overline{\xi})] - [H(\xi_i(x)) + \xi_i(x)\overline{F}(\xi_i(x))] - \overline{\xi}[\overline{F}(\overline{\xi}) - \overline{F}(\xi_i(x))]\\ &= \int_{\overline{\xi}}^\infty sf(s)ds - \int_{\xi_i(x)}^\infty sf(s)ds - \overline{\xi}[\overline{F}(\overline{\xi}) - \overline{F}(\xi_i(x))]\\ &= \int_{\overline{\xi}}^{\xi_i(x)} sf(s)ds - \overline{\xi}[\overline{F}(\overline{\xi}) - \overline{F}(\xi_i(x))]\\ &\geq \overline{\xi}\int_{\overline{\xi}}^{\xi_i(x)} f(s)ds - \overline{\xi}[\overline{F}(\overline{\xi}) - \overline{F}(\xi_i(x))] = 0.\end{aligned}$$

Thus, $\Delta > 0$, the desired contradiction. Hence, $\xi_i(x) < \overline{\xi}$. □

Clearly, $\xi_1(0) = \xi$, where ξ is given in (4). Because search ceases for $x \geq \overline{\xi}$ and $\xi_i(x) < \overline{\xi}$ for $x < \overline{\xi}$, we define $\xi_i(x) = x$ for $x \geq \overline{\xi}$.

THEOREM 2: When i implementations remain and the current state is x, the optimal return $V_i(x)$ is strictly increasing in i for $x < \bar{\xi}$.

PROOF: Fix $x < \overline{\xi}$ and $i \geq 1$. By Theorem 1, we know that $\xi_j(x) < \overline{\xi}$ for $j = 1, 2, \ldots, i$. Consequently, the fact that F is strictly increasing on an interval of the form $[a, \overline{\xi}]$ (see footnote 6) ensures $P(X_0 < \overline{\xi}) > 0$, where X_0 is the state when 0 periods remain and x is the state when i periods remain. Suppose $i + 1$ periods rather than i periods remain. Consider the policy which when $j + 1$ periods remain makes precisely the same decision that is made by an optimal policy when j periods remain, $j = i, i - 1, \ldots, 1$. If $X_0 \geq \overline{\xi}$, then this policy has the same return as the policy which is optimal when i implementations remain. But if $X_0 < \overline{\xi}$, which happens with positive probability, continue search until a technology of value v with $v > X_0$ is found. This induces a strict increase beyond the optimal return when i implementations remain, establishing $V_{i+1}(x) > V_i(x)$. □

4.2. Representing $V_i(x)$ as a Function of the Thresholds

The standard sequential search problem with search cost $c - x$ is precisely our problem in period 1 when the state is x. The opening of Section 4 presents the correspondence between the threshold $\xi_1(x)$ and the value of the standard sequential search problem: $V_1(x) = \xi_1(x)/\alpha$. Our generalization of this result, given in (14), represents V_i as the composition of $\xi_1, \ldots, \xi_i$ divided by α. If $x < \overline{\xi}$, then

$$V_i(x, y) = \left[x - c + \lambda \int_y^\infty V_{i-1}(s) f(s) ds\right] \Big/ (\alpha + \lambda \overline{F}(y)) \tag{11}$$

so that

$$\frac{\partial V_i(x, y)}{\partial y} = \frac{\lambda f(y)}{(\alpha + \lambda \overline{F}(y))^2} \left[-V_{i-1}(y)(\alpha + \lambda \overline{F}(y)) + x - c + \lambda \int_y^\infty V_{i-1}(s) f(s) ds\right].$$

Consequently, setting $dV_i(x, y)/dy = 0$, we see that $\xi_i(x)$ is the unique value of y that solves

$$V_{i-1}(y) = \left[x - c + \lambda \int_y^\infty V_{i-1}(s) f(s) ds\right] \Big/ (\alpha + \lambda \overline{F}(y)). \tag{12}$$

From (12) and(11), we conclude that

$$V_{i-1}(\xi_i(x)) = V_i(x, \xi_i(x)) = V_i(x). \tag{13}$$

COROLLARY 1: For $x < \bar{\xi}$, $\xi_i(x) > x$.

PROOF: If $\xi_i(x) = x$, then from (13) $V_i(x) = V_{i-1}(\xi_i(x)) = V_{i-1}(x)$, which contradicts Theorem 2. □

On reflection, the elegant equivalence between $V_i(x)$ and $V_{i-1}(\xi_i(x))$ is unsurprising. When i periods remain, the monopolist is indifferent between implementing a technological discovery of value $\xi_i(x)$ and using up one of his implementation opportunities versus continuing with his R&D efforts with all i opportunities remaining. Electing to continue R&D has value $V_i(x)$; electing to implement the discovery increases the state to $\xi_i(x)$ but diminishes the number of implementations by 1.

The recursive formula given in (13) enables us to express V_i in terms of $\xi_1, \ldots, \xi_i$. Because $V_0(x) = x/\alpha$, (13) asserts that $V_1(x) = V_0(\xi_1(x)) = \xi_1(x)/\alpha$. Substituting this expression into (13) produces $V_2(x) = V_1(\xi_2(x)) = \xi_1 \circ \xi_2(x)/\alpha$. Iterating this last result in (13) yields a simple relationship between the optimal return V_i and the i thresholds $\xi_1, \ldots, \xi_i$:

$$V_i(x) = \xi_1 \circ \xi_2 \circ \cdots \circ \xi_i(x)/\alpha. \tag{14}$$

We now utilize the characterization of the optimal return given in (14) to more elegantly verify $\xi_i(x) > x$ for $x < \overline{\xi}$. We know from Theorem 1 that $\xi_i(x) \geq x$. Suppose $\xi_i(x) = x$ for some pair i and x with $i \geq 1$ and $x < \overline{\xi}$. Employing (14), we obtain

$$\begin{aligned} V_i(x) &= \xi_1 \circ \xi_2 \circ \cdots \circ \xi_{i-1}(\xi_i(x))/\alpha \\ &= \xi_1 \circ \xi_2 \circ \cdots \circ \xi_{i-1}(x)/\alpha = V_{i-1}(x). \end{aligned}$$

But this contradicts Theorem 2.

4.3. Recursive Calculation of Thresholds

Next, we demonstrate that the threshold in each period is strictly increasing in the level x of the technology currently in place. Rearranging the first-order condition (12), with $i \geq 1$, we see that $\xi_i(x)$ is the unique value of y that solves

$$\alpha V_{i-1}(y) = x - c + \lambda \int_y^\infty [V_{i-1}(s) - V_{i-1}(y)] f(s) ds. \tag{15}$$

THEOREM 3: For each $i \geq 1$, the threshold $\xi_i(x)$ is a strictly increasing function of x for $x < \overline{\xi}$.

PROOF: Notice that the right-hand side of (15) increases in x. Because $V_{i-1}(y)$ strictly increases in y by Lemma 1, the right-hand side of (15) strictly decreases in y and the left-hand side of (15) strictly increases in y. Thus, an increase in x requires that $y = \xi_i(x)$ increase. □

The special case of $i = 1$ for Eq. (15) yields

$$\alpha V_0(y) = x - c + \lambda \int_y^\infty [V_0(s) - V_0(y)] f(s) ds. \quad (16)$$

Because $V_0(s) = s/\alpha$, (16) produces

$$y + c = x + \lambda H(y)/\alpha. \quad (17)$$

The y that solves (17) is $\xi_1(x)$. Of course, this is the same as (6).

Theorem 3 states that the higher the level of the installed technology, the higher the threshold. In particular, a small increase ϵ in the value x of the current technology when $i \geq 1$ periods remain increases the optimal return $V_i(x)$ by approximately $\epsilon/(\alpha + \lambda \overline{F}(\xi_i(x))$ because the return per unit time starting from state $x + \epsilon$ is ϵ larger for an exponential length of time with parameter $\lambda \overline{F}(\xi_i(x))$. This approximation is exact in the limit as $\epsilon \to 0$.

LEMMA 2: For $i \geq 1$, the derivative of $V_i(\cdot)$ satisfies

$$dV_i(x)/dx = 1/[\alpha + \lambda \overline{F}(\xi_i(x))], \quad \text{for } x < \overline{\xi}. \quad (18)$$

PROOF: Using $y^\star = \xi_i(x)$ in (11), we have $V_i(x) = V_i(x, y^\star)$ so

$$\frac{dV_i(x)}{dx} = \frac{\partial V_i(x, y^\star)}{\partial x} + \frac{\partial V_i(x, y^\star)}{\partial y^\star} \frac{\partial y^\star}{\partial x} = \frac{\partial V_i(x, y^\star)}{\partial x}$$
$$= 1/[\alpha + \lambda \overline{F}(y^\star)],$$

as $\partial V_i(x, y)/\partial y \,|_{y=\xi_i(x)} = 0$. □

We use (18) to demonstrate the relationship between the thresholds from one period to the next: for a given level of x, the threshold increases when fewer opportunities for implementation remain. The monopolist is more selective when fewer implementation opportunities remain.

THEOREM 4: For each $x < \overline{\xi}$, the optimal threshold $\xi_i(x)$ is a strictly decreasing function of i.

PROOF: We begin by dealing with a minor technicality. Because the set where the density f of F is positive is the interval (a, b) we can assume without loss of generality that $\xi_i(x) \geq a$. By footnote 6, $a < \overline{\xi} < b$, and $\xi_i(x) < \overline{\xi} \leq b$ by Theorem 1. Hence, if $\xi_i(x) < \xi_{i-1}(x)$, then $\overline{F}(\xi_i(x)) > \overline{F}(\xi_{i-1}(x))$.

The proof is by induction. Assume $\xi_i(x) < \xi_{i-1}(x)$ for all $x < \overline{\xi}$. In view of the formula for $V_i'(x)$ given in (18), the induction hypothesis implies

$$V_i'(x) < V_{i-1}'(x) \quad \text{for all } x < \overline{\xi}. \quad (19)$$

Consequently, (19) implies

$$\begin{aligned}\int_y^\infty [V_i(s) - V_i(y)] f(s) ds &= \int_y^\infty \left[\int_y^s V_i'(t) dt\right] f(s) ds \\ &< \int_y^\infty \left[\int_y^s V_{i-1}'(t) dt\right] f(s) ds \\ &= \int_y^\infty [V_{i-1}(s) - V_{i-1}(y)] f(s) ds.\end{aligned}$$

Thus, for each y, the right-hand side of (15) decreases in i. Also, $V_{i-1}(x)$ is strictly increasing in i for $x < \overline{\xi}$. Hence, $\xi_{i+1}(x) < \xi_i(x)$ for all $x < \overline{\xi}$. This verifies the induction hypothesis.

It only remains to show that $\xi_2(x) < \xi_1(x)$ for all $x < \overline{\xi}$. If there is an $x < \overline{\xi}$ such that $\xi_2(x) \geq \xi_1(x)$, then [immediately below we supply a reason for each comparison]

$$\begin{aligned}\xi_2(x) &< \xi_1(\xi_2(x)) = \alpha V_1(\xi_2(x)) \\ &= x - c + \lambda \int_{\xi_2(x)}^\infty [V_1(s) - V_1(\xi_2(x))] f(s) ds \\ &= x - c + \frac{\lambda}{\alpha} \int_{\xi_2(x)}^\infty \left[\int_{\xi_2(x)}^s \alpha V_1'(t) dt\right] f(s) ds \\ &< x - c + \frac{\lambda}{\alpha} \int_{\xi_2(x)}^\infty [s - \xi_2(x)] f(s) ds \\ &= x - c + \frac{\lambda}{\alpha} H(\xi_2(x)) \\ &= x - c + \frac{\lambda}{\alpha} H(\xi_1(x)) + \frac{\lambda}{\alpha}[H(\xi_2(x)) - H(\xi_1(x))] \\ &= \xi_1(x) + \frac{\lambda}{\alpha}[H(\xi_2(x)) - H(\xi_1(x))] \\ &\leq \xi_1(x),\end{aligned}$$

where we use, in order, $\xi_1(x)$ is strictly increasing (see Theorem 1); (14); (15); Fundamental Theorem of Calculus; $\alpha V_i'(x) < 1$ by (18) and $\overline{F}(\xi_i(x)) > \overline{F}(\overline{\xi}) > 0$; definition of H; adding and subtracting the same quantity; the first-order condition (17) for $\xi_1(x)$; and H strictly decreasing and the assumption $\xi_2(x) \geq \xi_1(x)$. This contradicts our assumption that $\xi_2(x) \geq \xi_1(x)$ for some $x < \overline{\xi}$. Hence, $\xi_2(x) < \xi_1(x)$ for all $x < \overline{\xi}$, thereby completing the induction argument. □

Theorems 1, 3, and 4 establish the structure of the optimal policy: the threshold increases in the incumbent technology and in the number of implementation opportunities remaining. Figure 2 gives a graphical representations of the optimal policy. Theorem 4 also establishes decreasing returns to additional opportunities to implement.

COROLLARY 2: If $x < \overline{\xi}$, then $V_{i+1}(x) - V_i(x) < V_i(x) - V_{i-1}(x)$.

PROOF: By (13) $V_i(x) = V_{i-1}(\xi_i(x))$. Hence,

$$\begin{aligned} V_{i+1}(x) - V_i(x) &= V_i(\xi_{i+1}(x)) - V_i(x) \\ &= \int_x^{\xi_{i+1}(x)} V_i'(u)du < \int_x^{\xi_{i+1}(x)} V_{i-1}'(u)du \\ &< \int_x^{\xi_i(x)} V_{i-1}'(u)du \\ &= V_{i-1}(\xi_i(x)) - V_{i-1}(x) \\ &= V_i(x) - V_{i-1}(x), \end{aligned}$$

where the second equality follows from $\xi_{i+1}(x) > x$; the first inequality follows from (19); the second inequality follows from $V' > 0$, Theorem 4, and $\xi(x) > x$; and the last two equalities follow from the Fundamental Theorem of Calculus and (13), respectively. □

Theorem 3 established that $\xi_i(x)$ is strictly increasing in x. We now augment this result and show that $\xi_i(x)$ increases more slowly than x.

Because $dH(x)/dx = -\overline{F}(x)$, differentiating the first-order condition (16) produces

$$\begin{aligned} \xi_1'(x) &= (\lambda/\alpha)H'(\xi_1(x))\xi_1'(x) + 1 \\ &= -(\lambda/\alpha)\overline{F}(\xi_1(x))\xi_1'(x) + 1, \end{aligned}$$

whence

$$\xi_1'(x) = \frac{1}{1 + (\lambda/\alpha)\overline{F}(\xi_1(x))}, \text{ for } 0 \le x < \overline{\xi}. \tag{20}$$

so that

$$0 < \xi_1'(x) < 1, \quad \text{for } 0 \le x < \overline{\xi}.$$

By (13), $V_i'(x) = \xi_i'(x)V_{i-1}'(\xi_i(x))$ so that

$$\begin{aligned} \xi_i'(x) &= \frac{V_i'(x)}{V_{i-1}(\xi_i(x))} \\ &= \frac{\alpha + \lambda\overline{F}(\xi_{i-1}(\xi_i(x))}{\alpha + \lambda\overline{F}(\xi_i(x))}, \quad \text{for } x < \overline{\xi} \quad \text{and } i \ge 2, \end{aligned}$$

where the last equality follows from (18). Theorem 1 asserts that $z \equiv \xi_{i-1}(\xi_i(x)) > \xi_i(x)$ so $\overline{F}(z) < \overline{F}(\xi_i(x))$. Hence, $\xi_i'(x) < 1$. In summary, we have

$$0 < \xi_i'(x) < 1 \quad \text{for } x < \overline{\xi} \quad \text{and } i \ge 1. \tag{21}$$

While the period i threshold is increasing in x, it is not increasing as fast as x: $\xi_i(x) - x$ is strictly positive but $\xi_i(x) - x \downarrow 0$ as $x \to \overline{\xi}$.

5. INCREASED ABILITY TO GENERATE NEW TECHNOLOGIES

The analysis in Section 4 addresses the monopolist's search for innovations in a stationary environment: the problem parameters, including the search cost c, the arrival rate λ of new technological discoveries, and the distribution F of the size of technological discoveries, are neither functions of time nor of the monopolist's decisions. There is, however, an enormous literature which unmistakably documents the marked impact of learning-by-doing as regards technological innovation. Just as practice leads to the improvement of most skills, the introduction of product improvements and the implementation of new technological discoveries leads to improvements in a firm's skill in generating new technological improvements. As noted by Mishina [27, p. 147], "experience is the mother of improvement."

In this regard, we are not referring to the learning curve[7] wherein the production cost per unit drops as the gross investment increases (see Arrow [2]) or as the cumulative number of units produced of a product increases (see Spence [33]).[8] Instead, we are referring to the phenomenon in which the firm's R&D ability improves when it steers an innovation through the complete cycle starting with discovery and ending with market introduction.

In this section, we examine how the monopolist's optimal decisions change when adoption of a new and better technology in and of itself leads to an increased ability to generate new technologies. This increased ability is modeled either as a decrease in the cost c of search, an increase in the arrival rate λ, or an improvement (in the sense of first-order stochastic dominance) in the distribution F. We do not alter the discount rate α. This improvement occurs naturally if the monopolist learns from the experience of bringing a product

[7] The learning curve is given by the relationship $y = ax^b$, where y_t is the input—cost or, in the case of the B-17 heavy bomber, direct labor hours—and x_t is cumulative output (cumulative number of bombers manufactured). As explained by Mishina (p. 147), "cumulative output makes sense as a proxy variable for some sort of production-related experience".
The Boston Consulting Group wrote on the subject in 1972, but Wright [35], director of engineering at the Curtiss-Wright Corporation, had already written cogently on the subject, in the context of building airframes.

[8] In 1990, Argote et al. [1] brought the famous case of the WW II Liberty Ships to the attention of Management Scientists. There has been an intense debate about whether or not the learning curve applies to the complexities of the Liberty Ships. But any potential for debate regarding the application of the learning curve to the B-17 heavy bomber (a.k.a. the Flying Fortress) is easily dispelled by Kazuhiro Mishina's fascinating article. He presents a compelling argument that, "It is the system of production that embodies learning, not the direct workers." Today, the study of the learning curve has progressed and, for example, Bayesian approaches to estimating the parameter b are considered (see Mazzola and McCardle [24]).

(or process improvement) to market. Learning-by-doing renders the discovery process faster (higher λ), better (improved F), or cheaper (lower c). See Lippman and McCardle [22].

We think of the examination of this increased generating ability as a well motivated comparative statics analysis. Though well motivated, effecting this particular comparative statics analysis turns out to be quite complex. Furthermore, in the process of carrying it out it became evident that extreme care must be exercised in applying the analysis of Section 4. Accordingly, we limited consideration to the case in which the parameters c, λ, and F change (i.e., improve) only in the very last period. As before, we say we are in period i when i innovation opportunities remain.

The caveats above apply even to this simple non-stationary (learning) environment. For example, it appears reasonable and innocuous to require that each new technology implemented improve on the currently installed technology. Innocuous as this requirement appears, it has strong implications: increased generating ability in period 1 can not only render recall crucial but also can scorch Theorem 1 by negating the optimality of thresholds (our fundamental structural result) and turn things upside down. We address these two issues in Examples 1 and 2. Example 1 provides a three period problem in which use of recall is optimal whereas Example 2 provides a three period problem in which the optimal policy lies in the class of reverse threshold rules. A reverse threshold rule is a policy which, given the state x and the period i with $i \geq 2$, implements a technological discovery if and only if it lies at or below a critical number $\eta_i(x)$. That is, with this seemingly innocuous requirement it can be sub-optimal to implement an innovation that is good and optimal to implement an innovation that is bad. The reason is simple: if the offer distribution F (or arrival rate λ) is fabulous in the last period (called period 1), then implementing a technological discovery with large value x when three periods remain can lead to a very long delay in entering the last period. [Notably, this issue never presents a problem in the penultimate period.]

EXAMPLE 1: Recall can be advantageous for the searcher.

Suppose that when $i = 1$ two offers are possible: either M or 0 with probabilities p_M and $1 - p_M$ respectively. When $i > 1$, the two possible values for the technology are m and 0 with probabilities p_m and $1 - p_m$ respectively, with $m << M$. Finally, suppose that $N = 3$, and the monopolist starts with no innovation implemented but having just received an innovation of value m. Should the monopolist implement the innovation now and move into period $i = 2$? Or should the monopolist set the current innovation aside, search again, accept (and implement) any arriving innovation, search again when $i = 2$, and implement the previously set aside implementation so as to enter period $i = 1$ after the first arrival in period $i = 2$ with an installed innovation of value m and thereby earn m while waiting for the high value (M) innovation in period $i = 1$?

If the monopolist decides to implement the current innovation, he will begin search in period $i = 2$, and he will implement the next arrival (without regard to its value). With probability $1 - p_m$ he will enter period $i = 1$ with an installed innovation of value 0 and with probability p_m of value m. His expected discounted return will be

$$\frac{10 - c}{\alpha + \lambda} + \frac{\lambda}{\alpha + \lambda}\left[\frac{mp_m - c}{\alpha + \lambda p_M} + \frac{\lambda p_M}{\alpha + \lambda p_m}\frac{M}{\alpha}\right].$$

On the other hand, if the monopolist does not implement the innovation, but immediately starts searching again (i remains at 3), he pays a cost c until the next innovation arrives. He implements that innovation (changing i to 2) and immediately begins search again. Upon the arrival of the very next innovation, he recalls the initial innovation of value m and implements it. He then enters the last period ($i = 1$) and continues searching until he encounters an innovation of value M. The expected return for this strategy is

$$\frac{-c}{\alpha + \lambda} + \frac{\lambda}{\alpha + \lambda} \times \left\{\frac{mp_m - c}{\alpha + \lambda} + \frac{\lambda}{\alpha + \lambda}\left[\frac{m - c}{\alpha + \lambda p_M} + \frac{\lambda p_M}{\alpha + \lambda p_M}\frac{M}{\alpha}\right]\right\}.$$

Set $m = 10$, $M = 1000$, $p_M = p_m = 1/1000$, $c = 0$, $\lambda = 10$, and $\alpha = 1.0$. In the former case, the discounted return is

$$= \frac{10}{11} + \frac{10}{11}\left[\frac{(1/100)}{1.01} + \frac{(1/100)}{1.01}1000\right] = 9.92.$$

In the latter case, the return is

$$= \frac{10}{11}\left[\frac{0.01}{11} + \frac{10}{11}\left[\frac{10}{1.01} + \frac{10}{1.01}\right]\right] = 16.37.$$

Of course, the monopolist could implement the innovation of value 10 when $i = 3$, search for another innovation of value 10 when $i = 2$, and then search for an innovation of value 1000 in the last period. A similar set of calculations reveals that the expected return from this strategy is 10.10. Because there are no other reasonable policies to consider, it is optimal to reject the innovation of value 10 when $i = 3$ and recall it after an arrival when $i = 2$. □

EXAMPLE 2: A reverse threshold rule can be optimal.

Throughout this example, we explicitly impose the restriction that a technological discovery can be implemented only if it is at least as good as the technology currently in place. We will show that with this explicit restriction in an environment with increased generating ability, it is possible for the optimal return function $V_i^+(x)$ to be strictly decreasing in x for $x < L$ where $L = inf\{x : F(x) = 1\}$ and $i \geq 2$. We begin with a remark.

REMARK: If $V_i^+(x)$ is strictly decreasing in x for $x < L$ and if $i \geq 2$, then the optimal policy in period $i + 1$ is a reverse threshold: there is a critical number $\eta_{i+1}(x)$ such that in period $i + 1$ it is optimal to implement a technology y if and only if $x \leq y \leq \eta_{i+1}(x)$.

In periods 1 and 2, the optimal policy is a threshold rule and $V_1^+(x)$ is strictly increasing in x. Fix $i \geq 2$ and assume $V_i^+(x)$ is decreasing in x for $x < L$. Let x_{i+1} be the state in period $i + 1$. If it is optimal to implement a technology of value $y > x_{i+1}$ in period $i + 1$ and if $y' < y$, then $V_i^+(y') > V_i^+(y) \geq V_{i+1}^+(x_{i+1})$, where the strict inequality follows the hypothesis that $V_i^+(s)$ is decreasing in s. Hence, it is optimal to implement y'. This completes the proof of the remark.

We now set $N = 3$ and exhibit the parameter values to produce an example in which a reverse threshold rule is optimal in period 3. To begin, set the offer distribution in period i to be a uniform random variable on (99,101) if $i = 1$ and a uniform random variable on $(0, 4)$ if $i > 1$. Set the discount rate, the arrival rate, and the search cost per unit time to 1 so $\lambda/(\alpha + \lambda) = 1/(\alpha + \lambda) = 1/2$. It is easily verified that starting from state 0, search is profitable in periods 1, 2, 3.

Of course, $V_0^+(x) = x/\alpha$. Because the state cannot exceed 4 at the beginning of period 1, in period 1 it is optimal to take the first technology to arrive. Consequently, $V_1^+(x) = (x-1)/2+(100/2)$. The return $V_2^+(x, y)$ when the state is x and threshold y is employed satisfies $V_2^+(x, y) = (y^2 + 198y - 16x - 792)/(4y - 32)$ which is decreasing in y for each $x \leq 4$ provided $y < 4$. If $y \geq 4$, then $V_2^+(x, y) = (x-c)/\alpha < x/\alpha$. In this case, it would be preferable to cease search so indeed it suffices to only consider $y < 4$. Hence, $\xi_2(x) = x$ and $V_2^+(x) = [x^2 + 182x - 792]/[4x - 32]$. For $0 < x < 4$, $V_2^+(x)$ is decreasing in x. Consequently, the Lemma above reveals that a reverse threshold policy is optimal in period 3. □

Henceforth, we prohibit recall and invoke a "use-it-or-lose-it" modeling mentality. Moreover, we drop the (implicit) requirement that a technological discovery can be implemented only if it is at least as good as the technology in place, and we think of the implementation of a technology inferior to the technology in place as an investment in learning. Just because an action is unprofitable in the short-run does not mean that it is an unwise action; in the long-run, it can engender future advantages.

The way out of this problem is to allow for the possibility that the learning value of implementing a less profitable discovery outweighs its immediate (short-run) economic loss. With this flexibility, the proof of Theorem 1 provides for the optimality of threshold policies.

Similar results (foregoing short-run profits for larger future gains) have been obtained by Cabral and Riordan [6] in a duopoly model with learning. In their model a firm can use the leverage offered by learning to engage in predatory pricing so as to drive competitors out of the market. Petrakis et al. [28] analyze a two period competitive economy with learning in the earlier period. They show that firms that choose not to exit will price below marginal cost in the earlier period.

The nonstationarity induced by increased generating ability introduces additional unanticipated difficulties. For example, there no longer exists a single threshold $\overline{\xi}$ which extinguishes search. Instead, for each period i there is a threshold $\overline{\xi}_i$ which extinguishes search in period i. This difficulty arises because the non-stationary problem is not a monotone stopping problem *à la* Chow and Robbins [8]: even though it does not pay to take precisely one more observation, it can pay to take several more observations. We can, however, show that these search extinguishing thresholds $\overline{\xi}_i$ are monotone (see Theorem 5). A second surprising and disturbing fact is that, contrary to Theorem 2, the return functions $V_i^+(x)$ are no longer increasing in i on $x < \overline{\xi}$. This result highlights the difficulties brought about by the confluence of three aspects of the model: costly search, discounting, and the dependence of the technology quality on experience.

In the standard (single adoption) search problem, the usual comparative statics analysis (in which one of these three parameters changes) is straightforward: a reduced cost of search, an increased arrival rate, or a better distribution (i.e., a stochastically larger distribution) induces an increase in the threshold (see Lippman and McCardle [22]). Of course, these results hold for period 1 because our one period problem starting in state x is simply the usual search problem with search cost equal to $c - x$. The improved generating ability unambiguously raises the thresholds; however, usually it is not possible to sign its impact upon the expected duration of search. This latter fact is important in our multi-period environment with learning.

The dynamic, cross-period effects, due to the impact of the monopolist's increased ability to generate improvements in period 1, are subtle. How does this increased generating ability in period 1 (the last period) change the optimal thresholds in earlier periods? This question is especially relevant for a firm that is investing in innovation capability.

Two opposing forces are at work. On the one hand, the period 1 increase in generating ability pushes the period i threshold down, $i \geq 2$: because the future (i.e., period 1) is more attractive than before, the period i threshold should be reduced to hasten the arrival of (the now more attractive future) period 1. On the other hand, because the period 1 thresholds increase, the duration of search in period 1 might increase. In this case, the monopolist must endure the flow rate of revenues associated with the technology implemented at the end of the penultimate period for a longer duration than in the stationary no-learning problem. This concern for "current" income might, and indeed often does, induce the monopolist to raise the period 2 threshold so as to reap an acceptable level of benefits during a longer duration of search in period 1. Theorem 6 shows that increased generating ability in period 1 does indeed induce a reduction in the thresholds in earlier periods; however, this result requires that the increased period 1 threshold not increase the expected duration of search in period 1. Absent a change in either F or λ, a decrease in c necessarily increases the threshold in period 1 and hence the duration of search in period 1. Consequently, the premise of Theorem 6 does not apply to a change in c when there is no change in λ and no change in F.

We denote the improved period 1 parameters by $c^+ \leq c$, $\lambda^+ \geq \lambda$, and F^+ stochastically larger than F (*i.e.* $\overline{F}^+(x) \geq \overline{F}(x)$). Of course, one of these three improvements will be strict. In the environment in which the period 1 parameters improve, we denote the optimal return and the optimal thresholds when i periods remain by appending a superscript $^+$ to yield $V_i^+(x)$ and $\xi_i^+(x)$.

Because of the nonstationarity, the level of technology which extinguishes search, label it $\overline{\xi}_i^+$, depends upon the period i: $\overline{\xi}_i^+ \equiv sup\{x : V_i^+(x) > x/\alpha\}$. Clearly, it behooves the monopolist to continue search if the current state x is less than $\overline{\xi}$. That is, $\overline{\xi}_i^+ \geq \overline{\xi}$ for $i \geq 1$.

THEOREM 5: The search extinguishing thresholds are non-increasing: $\overline{\xi} \leq \overline{\xi}_{i+1}^+ \leq \overline{\xi}_i^+$.

PROOF: We start by establishing that $V_i^+(x) > x/\alpha$ for all $x < \overline{\xi}_i^+$ whence search is optimal if $x < \overline{\xi}_i^+$. Fix $i > 1$ as our claim is clear from the analysis in Section 4 when $i = 1$. The proof used to establish (18) serves to establish $dV_i^+(x)/dx = 1/[\alpha + \lambda\overline{F}(\overline{\xi}_i^+(x))] \leq 1/\alpha$ Fix $x < \overline{\xi}_i^+$ and suppose $V_i^+(x) = x/\alpha$. There is a $y > x$ with $V_i^+(y) > y/\alpha$; otherwise, $\overline{\xi}_i^+ \leq x$. The mean value theorem asserts that there is a point z in (x, y) such that

$$1/\alpha \geq V_i^{+\prime}(z) = \frac{\left[V_i^+(y) - V_i^+(x)\right]}{(y-x)} > \frac{[y/\alpha - x/\alpha]}{(y-x)} = 1/\alpha,$$

a contradiction. Thus, it is optimal to search if and only if the state x is strictly less than $\overline{\xi}_i^+$.

Suppose $i \geq 2$ and $\overline{\xi}_i^+ < \overline{\xi}_{i+1}^+$ so that $V_{i+1}^+(\overline{\xi}_i^+) > \overline{\xi}_i^+/\alpha$ by the definition of $\overline{\xi}_{i+1}^+$. Set $z = \xi_{i+1}^+(\overline{\xi}_i^+)$. Either $z \geq \overline{\xi}_i^+$ or $z < \overline{\xi}_i^+$. We begin by assuming $z \geq \overline{\xi}_i^+$. We have

$$\begin{aligned}\overline{\xi}_i^+/\alpha < V_{i+1}^+(\overline{\xi}_i^+) &= \frac{\overline{\xi}_i^+ - c}{\alpha + \lambda\overline{F}(z)} + \frac{\lambda}{\alpha + \lambda\overline{F}(z)} \int_z^\infty V_i^+(t) f(t)dt \\ &= \frac{\overline{\xi}_i^+ - c}{\alpha + \lambda\overline{F}(z)} + \frac{\lambda}{\alpha + \lambda\overline{F}(z)} \int_z^\infty \frac{t}{\alpha} f(t)dt \\ &= \left\{\overline{\xi}_i^+ - c + \frac{\lambda}{\alpha}H(z) + \frac{\lambda}{\alpha}z\overline{F}(z)\right\} \Big/ [\alpha + \lambda\overline{F}(z)].\end{aligned}$$

Using (9) to substitute $\frac{\lambda}{\alpha}H(\overline{\xi})$ for c, multiplying the above inequality by $\frac{\alpha}{\lambda}[\alpha + \lambda\overline{F}(z)]$, and simplifying, we obtain

$$\begin{aligned}0 &< -H(\overline{\xi}) - \overline{\xi}_i^+\overline{F}(z) + H(z) + z\overline{F}(z) \\ &= -\int_{\overline{\xi}}^z \overline{F}(t)dt + \overline{F}(z)\left(z - \overline{\xi}_i^+\right) \leq 0\end{aligned}$$

because $\int_{\overline{\xi}}^z \overline{F}(t)dt \geq \int_{\overline{\xi}}^z \overline{F}(z)dt = \overline{F}(z)(z - \overline{\xi})$. This contradiction shows that $z \geq \overline{\xi}_i^+$ is not possible when $\overline{\xi}_i^+ < \overline{\xi}_{i+1}^+$.

Next, assume that $z < \overline{\xi}_i^+$. In this case $V_i^+(t) < V_i^+(\overline{\xi}_i^+) = \overline{\xi}_i^+/\alpha$ for $z < t < \overline{\xi}_i^+$ whence

$$\begin{aligned}\int_z^\infty V_i^+(t) f(t)dt &< \int_z^{\overline{\xi}_i^+} \frac{\overline{\xi}_i^+}{\alpha} f(t)dt + \int_{\overline{\xi}_i^+}^\infty \frac{t}{\alpha} f(t)dt \\ &= \frac{\overline{\xi}_i^+}{\alpha}\left(F(\overline{\xi}_i^+) - F(z)\right) + H(\overline{\xi}_i^+)/\alpha + \frac{\overline{\xi}_i^+}{\alpha}\overline{F}(\overline{\xi}_i^+) \\ &= \frac{\overline{\xi}_i^+}{\alpha}\overline{F}(z) + H(\overline{\xi}_i^+)/\alpha.\end{aligned}$$

Continuing as in the case $z \geq \overline{\xi}_i^+$, we use the inequality directly above, (9), and H non-increasing to produce

$$\begin{aligned}&\overline{\xi}_i^+/\alpha < V_{i+1}^+(\overline{\xi}_i^+) \\ &\quad < \frac{\overline{\xi}_i^+ - c}{\alpha + \lambda\overline{F}(z)} + \frac{\lambda}{\alpha + \lambda\overline{F}(z)}\left[\frac{\overline{\xi}_i^+}{\alpha}\overline{F}(z) + H(\overline{\xi}_i^+)/\alpha\right] \\ &\quad = \overline{\xi}_i^+/\alpha + \frac{\lambda}{\alpha}\frac{H(\overline{\xi}_i^+) - H(\overline{\xi})}{\alpha + \lambda\overline{F}(z)} \leq \overline{\xi}_i^+/\alpha,\end{aligned}$$

whence this contradiction shows that $z < \overline{\xi}_i^+$ is not possible when $\overline{\xi}_i^+ < \overline{\xi}_{i+1}^+$. We have now established that $\overline{\xi}_{i+1}^+ \leq \overline{\xi}_i^+$ for $i \geq 2$. With one minor exception, the proof for the case $i = 1$ is the same as above: the parameters become λ^+, c^+, and F^+ and necessarily $\xi_1^+ > \overline{\xi}$. □

In fact, we can obtain a more general comparative statics result. The proof mimics the proof of Theorem 5.

COROLLARY 3: If $\lambda_{i+1} \le \lambda_i, c_{i+1} \ge c_i$, and $\overline{F}_i(t) \ge \overline{F}_{i+1}(t)$ for all t and $i = 1, \ldots, N-1$, then $\overline{\xi} \le \overline{\xi}^+_{i+1} \le \overline{\xi}^+_i, i = 1, \ldots, N$.

THEOREM 6: Because of increased generating ability in period 1, $\xi_1^+(x) > \xi_1(x)$ for $x < \overline{\xi}$. If increased generating ability in period 1 does not increase the expected duration of search in period 1, then $\xi_i^+(x) < \xi_i(x)$ for $i \ge 2$ and $x < \overline{\xi}$.

PROOF: The proof is similar to the proof of Theorem 4. From (11), we see that the parameters in (15) match those of the threshold being obtained: when $\alpha V^+_{i-1}(x)$ is on the left-hand side of (15), the arrival rate and offer distribution are those of period i. Thus, for $i \ge 2$, the appropriate arrival rate and offer distribution are λ and F. When $i = 1$, (15) reduces to (16) so that these parameters become λ^+ and F^+ in the $^+$ environment.

For $i = 1$, the improved generating ability strictly increases the optimal return whence $\xi_1^+(x)/\alpha = V_1^+(x) > V_1(x) = \xi_1(x)/\alpha$.

The duration of search in period 1 is an exponential random variable with parameter $\lambda^+\overline{F}^+(\xi_1^+(x))$ or $\lambda\overline{F}(\xi_1(x))$, whence the expected duration of search is the inverse of these quantities. By hypothesis, in period 1 the expected duration of search in the $^+$ environment is not greater than that in the no-learning environment. That is, $\lambda^+\overline{F}^+(\xi_1^+(x)) \ge \lambda\overline{F}(\xi_1(x))$. Consequently, it follows from Lemma 1 that $dV_1^+(x)/dx \le dV_1(x)/dx$ so the right-hand side of (15) in the $^+$ environment is less than or equal to the right-hand side of (15) in the no-learning environment. Using the fact that $V_2^+(x)$ is strictly greater than $V_2(x)$ for all $x < \overline{\xi}$, it follows as per the proof of Theorem 4 that $\xi_2^+(x) < \xi_2(x)$, as desired.

Assume that $\xi_i^+(x) < \xi_i(x)$ for some i. Coupling this assumption with Lemma 1 and the fact that the arrival rate and offer distribution in period i are λ and F in both environments, it follows that $dV_i^+(x)/dx < dV_i(x)/dx$. The induction argument now follows as per the proof of Theorem 4. □

Examples 3 and 4 below treat the case of a stochastically larger distribution F and an increase in the arrival rate λ, respectively. These examples supply "general" conditions under which the increased generating ability in period 1 leads to an increase in the acceptance rate of new arrivals and hence a decrease in the expected duration of search in period 1. That is, these examples satisfy the hypothesis of Theorem 5.

EXAMPLE 3: Reduction in the expected duration of search due to a shift in F.

Suppose $X^+ = X + \delta$ so that $F^+(t) = F(t-\delta) \le F(t)$ with $\delta > 0$. We claim that

$$\xi_1(x) < \xi_1^+(x) < \xi_1(x) + \delta.$$

Note that the period 1 thresholds $\xi_1(x)$ and $\xi_1^+(x)$ in the no-learning environment and the $^+$ environment solve (16). Also note that the expected gain $H^+(x)$ satisfies $H^+(y) \equiv \int_y^\infty (s-y) f^+(s)ds = H(y-\delta)$ so that $H^+(y) > H(y)$ because H is strictly decreasing (on the interval where f is positive).

If $\xi_1^+(x) \ge \xi_1(x) + \delta$, then by (17)

$$\begin{aligned} x + \frac{\lambda}{\alpha} H(\xi_1(x)) - \xi_1(x) = c &= x + \frac{\lambda}{\alpha} H^+\big(\xi_1^+(x)\big) - \xi_1^+(x) \\ &= x + \frac{\lambda}{\alpha} H\big(\xi_1^+(x) - \delta\big) - \xi_1^+(x) \\ &\le x + \frac{\lambda}{\alpha} H(\xi_1(x)) - \xi_1^+(x) \\ &< x + \frac{\lambda}{\alpha} H(\xi_1(x)) - \xi_1(x), \end{aligned}$$

where the two inequalities follow from H decreasing and $\xi_1^+(x) \ge \xi_1(x) + \delta$ and from $\xi_1^+(x) > \xi_1(x)$, a contradiction. Thus, $\overline{F}^+(\xi_1^+(x)) = \overline{F}(\xi_1^+(x) - \delta) > \overline{F}(\xi_1(x))$. Hence, the expected duration of search in period 1 is strictly less in the learning environment. □

EXAMPLE 4: Reduction in the expected duration of search due to an increase in λ.

Suppose that $\lambda^+ > \lambda_0$ where λ^+ and λ_0 are the period 1 arrival rates in the $^+$ and the no-learning environments. Further, suppose that the offer distribution F has hazard rate function h that is non-decreasing. Of course, $h(t) = f(t)/\overline{F}(t)$: both the exponential and uniform distributions have a non-decreasing hazard rate function.

Let y_λ denote the optimal threshold in the one period problem when the arrival rate is λ. Thus, y_λ is the unique solution to (16): $y_\lambda = x - c + \frac{\lambda}{\alpha} H(y_\lambda)$. Differentiating (16) with respect to λ produces $y'_\lambda = H(y_\lambda)/[\alpha + \lambda\overline{F}(y_\lambda)] < H(y_\lambda)/\lambda\overline{F}(y_\lambda)$. Using the well-known equalities $\overline{F}(t) = exp\{-\int_0^t h(s)ds\}$ and $H(t) = \int_s^\infty \overline{F}(s)ds$, the inequality above yields

$$\begin{aligned} 0 < y'_\lambda &< \int_{y_\lambda}^\infty \overline{F}(s)ds/[\lambda\overline{F}(y_\lambda)] \\ &= \frac{1}{\lambda}\int_{y_\lambda}^\infty e^{-\int_{y_\lambda}^s h(t)dt} ds < \frac{1}{\lambda}\int_{y_\lambda}^\infty e^{-(s-y_\lambda)h(y_\lambda)} ds \\ &= \frac{1}{\lambda h(y_\lambda)}, \end{aligned}$$

whence

$$y_{\lambda^+} - y_{\lambda_0} = \int_{\lambda_0}^{\lambda^+} y'_\lambda d\lambda < \int_{\lambda_0}^{\lambda^+} \frac{1}{\lambda h(y_\lambda)} d\lambda \leq \frac{1}{h(y_{\lambda_0})} \cdot \ln(\lambda^+/\lambda_0), \quad (22)$$

where the last inequality follows from y_λ increasing and h non-decreasing. Use of (22) yields

$$\frac{d\lambda \overline{F}(y_\lambda)}{d\lambda} = \overline{F}(y_\lambda) - \lambda f(y_\lambda) y'_\lambda > \overline{F}(y_\lambda) - \frac{\lambda f(y_\lambda)}{\lambda f(y_\lambda)/\overline{F}(y_\lambda)} = 0.$$

This verifies that the period 1 expected duration of search in the increased generating ability environment is strictly less than in the no-learning environment. □

As shown in Examples 3 and 4, the increased generating ability in period 1 reduces the expected duration of search in period 1 whence the hypothesis of Theorem 6 is satisfied. Some improvements in the monopolist's generating ability in period 1, however, do not reduce the expected duration of search in period 1. For example, a reduction in the period 1 search cost necessarily increases the expected duration of search. Furthermore, a period 1 improvement that increases the expected duration of search might induce an increase in thresholds for period i with $i \geq 2$. Example 5 demonstrates that this possibility can occur.

EXAMPLE 5: Reducing the search cost in period 1 induces an increase in a period 2 threshold For simplicity, the offer distribution F is discrete. Set $c_1 = c_2 = 1.5$ and $c_1^+ = 0.5$, and also set $\alpha = \lambda = 1$, $P(X = 1) = .1$, $P(X = 2) = .8$, and $P(X = 7.99) = .1$. Necessarily, the reduction in the period 1 search cost increases the expected duration of search in period 1.

With $c_1 = 1.5$, search is profitable in period 1 only when $x = 0$. This holds because when $x = 1$, it is better to stop (and earn 1 per unit time) than to search (and, with a threshold of 2, earn .99947). The same conclusions hold for period 2: search only if $x = 0$.

When $c_1^+ = 0.5$, search in period 1 is profitable provided there are gains to be made (i.e., for any level of $x < 7.99$). In period 2, search is profitable for $x = 1$, and the best threshold is 2. Therefore, $\xi_2^+(1) = 2 > 1 = \xi_2(1)$. □

These examples, together with the results presented earlier in this section, show how changes in period 1 parameters have major consequences for both the optimal policy and the nature of the return functions V_i^+ for $i > 1$. In particular, unlike V_i, the return functions V_i^+ need not be increasing in i. Examples 3 and 4 reveal that better (stochastically larger offer distribution F) and faster (larger λ), respectively, can reduce the duration of search and therefore reduce the thresholds at earlier times. On the other hand, as per Example 5, a reduction in the period 1 search cost can lead to an increase in the thresholds at earlier times. Thus, either of the two forces, an increased incentive to get to period 1 quickly and utilize the improved generating ability or an increased concern for the magnitude of the state of the technology in place in period 1, can dominate.

6. CONCLUSION

Our contribution has been to model and analyze R&D efforts in a search setting in which multiple discoveries can be sequentially implemented and, quite importantly, search becomes less rewarding as more discoveries are implemented. Other work in the literature has allowed for sequential implementations, but it assumes that R&D effort necessarily leads to improvement while our model captures the realities of innovation refinement as discussed by economic historians for many decades.

As in other search models, the optimal policy entails thresholds. At any point in time there are two thresholds of interest: if the next discovery exceeds the lower but not the higher threshold, it is optimal to implement it and to continue search; if the next discovery exceeds the higher threshold, then it is optimal to implement it and to stop searching entirely. While the lower of the two thresholds depends on both the currently implemented technology and the maximum number of implementations remaining, the higher threshold, that which extinguishes search, is independent of these two factors. Moreover, as the number i of innovations that can be implemented increases, the threshold $\xi_i(x)$ decreases. Thus, the flexibility to introduce more versions of the product (stochastically) reduces the average size of each improvement brought to market. In this vein, we anticipate witnessing many small innovations when the innovations require little cognitive effort on the part of customers. For example, updates to Microsoft's Windows operating system, which can be downloaded from the Internet, are frequent (sometimes weekly), and they constitute small changes. Because adoption of these changes is nearly painless, consumers are willing to tolerate many such small introductions (so N is large), and each introduction is minor. In contrast, entirely new versions of the operating system require considerably more user involvement (including installing the new software and dealing with the resultant compatibility issues). Because these major introductions entail substantial user investment of time and energy, they engender customer resistance. Consequently, N is small, introductions are less frequent, and the magnitude of the change from one introduction to the next is much larger.

However, the description of our model has been in terms of a firm's repeated introductions of product improvements, the model and its results apply more broadly. As discussed in Lippman and McCall [21], search should be viewed as a broad paradigm, a flexible conceptual framework that applies equally well to job search and R&D: "search has been effectively employed in modeling directly productive activities …. The productive efforts necessary to bring about technological change are often modeled mathematically as if the effort were search activity." Our model applies when the crucial feature of the situation under study entails the repeated decision to act now or delay. For example, a proposed but imperfect resolution to a dispute can be acceptable to both parties as "good enough for now" or preferable to "continue to negotiate and expend more resources." The current resolution may be adopted even as the parties continue to negotiate and search for a better one. This dispute resolution example illustrates how our analysis can be applied in contexts far beyond technology: it applies not merely to some new thing (i.e., technology) but also to any new way of doing something. Our arguments apply to situations in which new alternatives are yet to be generated and interim solutions can be implemented, like climbing the rungs of a ladder.

A crucial feature of our model is the fact that N, the maximum number of new discoveries that can be implemented, is finite. If $N = \infty$, then it is optimal to implement each and every new discovery that improves on the current technology and search ceases only upon finding a technology of value $\overline{\xi}$ or better. But the justification for $N < \infty$ is not limited to mathematical necessity. Practical considerations also suggest $N < \infty$. Each new implementation not only further saturates the market but also induces consumers not to purchase for at least two reasons. First, in the spirit of the "durable goods monopoly" problem (see Coase [10] and Bulow [4]), the new implementation suggests to consumers that further improvements are just around the corner. Why purchase now when a purchase can be made at a lower price later.[9] Second, just as too many cooks spoil the broth, a multitude of product changes annoys consumers: after a small number of new product introductions, consumers are resistant to change. Having learned how to use some new product, such as software, the advantage bestowed by a bevy of new features can be outweighed by the cost of learning how to use them: change exerts a toll. Additionally, many product modifications or upgrades induce consumer confusion. Rather than adopt each new innovation, a consumer (and employee too) might elect to skip a few steps on the ladder of technological improvement and later leap-frog onto a higher rung. Finally, frequent modifications occlude the customer's insight and impede progress along the learning curve.

Section 5 formally addressed the issue of learning from experience. If an implementation provides lessons that improve the search environment—making discoveries arrive faster, better, or cheaper—then it can be optimal for the firm to relax it's standards (i.e., lower the threshold) to learn the lessons. However, we have shown that the opposite effect is possible too: a future improvement in the search environment can cause the firm to raise its current standards (if that future improvement serves to increase the duration of search).

There are several extensions of this work that we have not covered. Lippman and Mamer [18] model competitive search for a single introduction. What happens to the decision rules in an environment with competition and interim solutions? Another avenue worth exploring is other effects of delay, such as a deterioration of the profit potential of the market.

[9] Once the high-valuation customers have purchased, it is optimal for the monopolist to reduce the price; anticipating this price reduction prompts high-valuation customers to postpone their purchase.

REFERENCES

[1] L. Argote, S.L. Beckman, and D. Epple, The persistence and transfer of learning in industrial settings, Manage Sci 36 (1990), 140–154.
[2] K.J. Arrow, The economic implications of learning by doing, Rev Econ Studies 29 (1962), 155–173.
[3] Y. Balcer and S.A. Lippman, Technological expectations and adoption of improved technology, J Econ Theory 34 (1984), 292–318.
[4] J.I. Bulow, Durable-goods monopolists, J Polit Econ 90 (1982), 314–332.
[5] K. Burdett, A theory of employee job search and quit rates, Am Econ Rev 68 (1978), 212–220.
[6] L. Cabral and M. Riordan, The learning curve, market dominance, and predatory pricing, Econometrica 62 (1994), 1115–1140.
[7] F.G. Cauley and S.A. Lippman, Myopia and R&D/production complementarities, Econ Theory 4 (1994), 437–451.
[8] Y.S. Chow and H. Robbins, "A martingale system theorem and applications," in: Proceedings Fourth Berkeley Symposium on Math. Statist. and Prob., Vol. 1, University of California Press, Berkeley and Los Angeles, 1961, pp. 93–104.
[9] Y.S. Chow, H. Robbins, and D. Siegmund, Great expectations: The theory of optimal stopping, Houghton Mifflin Company, Boston, 1971.
[10] R.H. Coase, Durability and monopoly, J Law Econ 15 (1972), 143–149.
[11] W.M. Cohen and D.A. Levinthal, Absorptive capacity: A new perspective on learning and innovation, Admin Sci Q 35 (1990), 128–152.
[12] W.M. Cohen and D.A. Levinthal, Fortune favors the prepared firm, Manage Sci 40 (1994), 227–251.
[13] C. Derman and J. Sacks, Replacement of periodically inspected equipment (An optimal stopping rule), Nav Res Logist Q 7 (1960), 597–607.
[14] U. Doraszelski, The net present value method versus the option value of waiting: A note on Farzin, Huisman and Kort (1998), J Econ Dyn Control 25 (2001), 1109–1115.
[15] Y.H. Farzin, K.J.M. Huisman, and P.M. Kort, Optimal timing of technology adoption, J Econ Dyn Control 22 (1998), 779–799.

[16] R.E. Gomory, From the 'Ladder of Science' to the product development cycle, Harvard Business Rev 67 (1989), 99–105.
[17] L.J. Kornish, On optimal replacement thresholds with technological expectations, J Econ Theory 89 (1999), 261–266.
[18] S.A. Lippman and J.W. Mamer, Preemptive innovation, J Econ Theory 61 (1993), 104–119.
[19] S.A. Lippman and J.J. McCall, The economics of job search: A survey, Econ Inquiry 14 (1976), 155–189.
[20] S.A. Lippman and J.J. McCall, "The economics of uncertainty: Selected topics and probabilistic methods," Chapter 6, in: Handbook of Mathematical Economics, K.J. Arrow and M.J. Intriligator (Editors), North-Holland Publishing Company, Amsterdam, 1981, pp. 211–284.
[21] S.A. Lippman and J.J. McCall, Search and the development of the economics of information, Estudios de Economia 20 (1993), 225–249.
[22] S.A. Lippman and K. McCardle, Does cheaper, faster, or better imply sooner in the timing of innovation decisions?, Manage Sci 33 (1992), 1058–1064.
[23] E. Mansfield, The economics of technological change, W.W. Norton & Company, Inc., New York, 1968.
[24] J.B. Mazzola and K.F. McCardle, A Bayesian approach to managing learning-curve uncertainty, Manage Sci 42 (1996), 680–692.
[25] R.A. Miller and K. Voltaire, A sequential stochastic tree problem, Econ Lett 5 (1980), 135–140.
[26] R.A. Miller and K. Voltaire, A stochastic analysis of the tree paradigm, J Econ Dynam Control 6 (1983), 371–386.
[27] K. Mishina, "Learning by new experiences: Revisiting the Flying Fortress Learning Curve," Learning by Doing in Markets, Firms, and Countries, N.R. Lamoreaux and D. M. G. Raff and P. Temin (Editors), University of Chicago Press, 1999, pp. 145–179.
[28] E. Petrakis, E. Rasumussen, and S. Roy, The learning curve in a competitive industry, RAND J Econ 28 (1997), 249–269.
[29] N. Rosenberg, Factors affecting the diffusion of technology, Explor Econ Hist 9 (1972), 3–33.
[30] S.M. Ross, Applied probability models with optimization applications, Dover Publications, Inc., New York, 1970.
[31] F.M. Scherer, Innovation and growth, The MIT Press, Cambridge, Massachusetts, 1984.
[32] J.A. Schumpeter, Capitalism, socialism, and democracy, Harper & Brothers, New York, 1942.
[33] A. Michael Spence, The learning curve and competition, Bell J Econ 12 (1981), 49–70.
[34] A.P. Usher, A history of mechanical inventions, Harvard University Press, Cambridge, Massachusetts, 1954.
[35] T.P. Wright, Factors affecting the cost of airplanes, J Aeronaut Sci 3 (1936), 122–128.