

ECONOMIC RESEARCH REPORTS

***LEARNING BY DOING AND
THE CHOICE OF TECHNOLOGY***

by **Boyan Jovanovic**
and
Yaw Nyarko

RR # 96-25

August 1996

**C.V. STARR CENTER
FOR APPLIED ECONOMICS**



**NEW YORK UNIVERSITY
FACULTY OF ARTS AND SCIENCE
DEPARTMENT OF ECONOMICS
WASHINGTON SQUARE
NEW YORK, NY 10003-6687**

Learning by Doing and the Choice of Technology¹

by

Boyan Jovanovic[†]

and

Yaw Nyarko[†]

August 4, 1996

Abstract

This is a one-agent model of learning by doing and technology choice. The more the agent uses a technology, the better he learns its parameters, and the more productive he gets. This expertise is a form of human capital.

Any given technology has bounded productivity, which therefore can grow in the long run only if the agent keeps switching to better technologies. But a switch of technologies temporarily reduces expertise: The bigger is the technological leap, the bigger the loss in expertise. The prospect of a productivity drop may prevent the agent from climbing the technological ladder as quickly as he might. Indeed, an agent may be so skilled at some technology that he will never switch again, so that he will experience no long run growth. In contrast, someone who is less skilled (and therefore less productive) at that technology may find it optimal to switch technologies over and over again, and therefore enjoy long-run growth in output. Thus the model can give rise to overtaking.

Keywords: Human Capital, growth, overtaking.

Journal of Economic Literature No's: C11, D8, O12, O3

[†] Department of Economics, New York University

¹We thank the C.V. Starr Center for Applied Economics at New York University for technical and financial help, and Chung Tse for help with the research.

1. Introduction

This paper explores a one-agent Bayesian model of learning-by-doing and technological choice. In the model, experience yields information, which improves decisions and raises productivity. Once the productivity gains on a given technology are exhausted, further growth can occur only by switching to a better technology. How transferable is the previously acquired knowledge in the new activity? This will depend on how similar the new activity is to the old, and this depends on how correlated their unknown parameters are. In this sense, transferable information is general human capital, and nontransferable information is specific human capital.

In its focus on the choice between sticking to a "current" technology and switching to a better one, our model is closest in spirit to that of [10]. Our paper adds to the literature by analyzing the full dynamics. Parente, as well as [4], had looked only at constant growth paths, and did not ask whether these were reachable from arbitrary initial conditions. Globally stable long-run growth can arise in our model too: For some values of the parameters, the agent switches to a new technological grade infinitely often. But for other parameter values, an agent will stick to an old technology for ever, and experience long run stagnation. Moreover, the long-run value of the growth rate may depend on initial conditions. In particular, "overtaking" may occur -- an agent may be so skilled at a technology that he will refuse to switch to a better, but unfamiliar one. Such an agent may in the long run be overtaken by an agent who initially is less skilled and less attached to the technology at hand, and who therefore is more willing to try a new one.

2. Technological Deepening and Opening

We now model first technological deepening, by which we mean learning more about a given grade of a technology, and technological opening, by which we mean learning about a new technology.

2.1. *Learning a Technology*

A risk-neutral agent can produce a good with one of several grades of a technology indexed by $n \in [0, \infty)$. If he uses grade n at date t , a decision z yields net output via the production function²

$$(1a) \quad q = \gamma^n [1 - (y_{nt} - z)^2], \quad \gamma > 1$$

where

$$(1b) \quad y_{nt} = \theta_n + w_{nt}$$

is a random target that fluctuates around a grade-specific parameter θ_n , and where w_{nt} is an i.i.d. normal variate with mean zero and variance σ_w^2 . The agent knows γ and knows the distribution of w_{nt} . The agent does not know θ_n but has some prior beliefs about it. Let $E_t(\cdot)$ denote the conditional expectation at t , and $\text{Var}_t(\cdot)$ the conditional variance. Then the optimal decision and the resulting expected output are³

$$(2a) \quad z = E_t(y_{nt}) = E_t(\theta_n), \quad \text{and}$$

$$(2b) \quad E_t(q) = \gamma^n [1 - \text{Var}_t(\theta_n) - \sigma_w^2].$$

Equation (2b) allows us to think of the posterior precision on θ_n as an index of human capital: If the agent uses grade n , he also observes y_{nt} and learns more about θ_n , which allows him to make a better decision z . This reduces the posterior variance, $\text{Var}_t(\theta_n)$, and raises his expected net output. The learning process is bounded -- using grade n of a technology forever allows the agent to learn θ_n completely so that $E_t(q) \rightarrow \gamma^n [1 - \sigma_w^2]$, which is finite for fixed n .

2.2. The Transfer of Human Capital across Grades of Technology

There is no direct cost of switching to a different grade of technology, and no cost to adjusting z . The link between grades is informational: We suppose that the relation between θ_n and θ_{n+k} for any n and $k \geq 0$ is:

$$(3) \quad \theta_{n+k} = \alpha^{k/2} \theta_n + e_k \quad \text{where } e_k \sim N(0, \rho_k \sigma_\epsilon^2) \quad \text{and} \quad \rho_k = \begin{cases} (1-\alpha^k)/(1-\alpha) & \text{when } \alpha \neq 1 \\ k & \text{when } \alpha = 1 \end{cases}$$

and where θ_n and e_k are independent. (Eq. (3) generalizes the AR-1 process for $k=1$, $\theta_{n+1} = \sqrt{\alpha} \theta_n + \epsilon_{n+1}$, to a diffusion process). A feature of (3) is that if $\alpha = 1$ and $\alpha_\epsilon^2 = 0$, then $\theta_{n+k} = \theta_n$ for all k , so that human capital is general and freely transferable across grades of technology. If $\alpha = 0$, human capital is grade-specific.

We assume that the prior over θ_1 at date 1 is normal. eq. (3), and the normality of w_{nt} imply that the posterior belief at each date over the parameter of any grade, θ_n , will also be normal. We define the following functions of x and portray them in fig. 1:

$$\begin{aligned}
h_1(x) &\equiv \sigma_w^2 x / (\sigma_w^2 + x); && \text{(updating)} \\
h_2(x,k) &\equiv \alpha^k x + \rho_k \sigma_\epsilon^2; && \text{(transfer of knowledge)} \\
h(x,k) &\equiv h_1(h_2(x,k)). && \text{(transfer followed by updating)}
\end{aligned}$$

Figure 1 around here.

These functions have the following interpretations: Suppose that grade n is the highest grade that the agent has worked with at the end of date t and suppose his posterior distribution over θ_n is normal with variance $x_{n,t} = \text{var}_t(\theta_n)$. If he uses grade n at date $t+1$, he will then see $y_{n,t+1}$, after which his posterior variance over grade n becomes, via Bayesian updating, $h_1(x_{n,t})$. If instead at date $t+1$ he chooses grade $n+k$, $h_2(x_{n,t},k)$ is his prior variance at date t over θ_{n+k} ; he will then see $y_{n+1,t+k}$, and his posterior variance over θ_{n+k} will be $h(x_{n,t},k)$.

Let x_k^{**} be the fixed point of the $h(\cdot,k)$ map. Since $h(0,k) > 0$ and h is continuous, bounded and concave (see Figure 1), x_k^{**} exists and is unique. Suppose that at each date the agent chooses a jump size of k . Then the posterior variance over the most recent grade chosen at date t , x_t , will be given by the t -th iterate of the $h(\cdot,k)$ mapping from $x=x_1$. The sequence $\{x_t\}_{t=1}^\infty$ converges monotonically to x_k^{**} .

Throughout we shall assume that the agent is myopic: The agent maximizes current period return in each period. Define $G(x,k)$ to be the expected net output from initial posterior variance x when a jump of size k is chosen and the status quo grade is grade 0. Then

$$(4) \quad G(x,k) = \gamma^k [1 - \sigma_w^2 - \alpha^k x - \rho_k \sigma_\epsilon^2] .$$

Define x^* to be the value of x such that $G(x,0) = G(x,1)$. Since $G(\cdot,0)$ and $G(\cdot,1)$ are linear in x with slopes 1 and $\alpha\gamma$ respectively, x^* is well-defined whenever $\alpha\gamma \neq 1$ (which, for ease of exposition, we assume throughout). Throughout, we impose a "No Recall" constraint: Once a grade has been passed over for a higher grade, it is never recalled. Moreover, we constrain $k \leq 1$: any jump size k in $[0,1]$ is feasible.

3. Case A (Overtaking; Beliefs affect long run growth catastrophically)

We define Case A by the following requirements:

- i. $\alpha\gamma < 1$; ii. $G(0,0) > G(0,1)$ and⁴ iii. $x^* < x_1^{**}$.

Even though any jump size k in $[0,1]$ is feasible, Proposition 3.1 will state that in Case A, only $k=0$ or $k=1$ will be chosen. This permits the representation in fig. 2, which shows the expected payoff from choosing $k=0$ (the action "STICK") and from $k=1$ (the action "SWITCH").

Figure 2 around here.

If the initial belief has variance x less than x^* , fig. 2 shows that the action STICK dominates SWITCH. Hence when $x < x^*$ the agent will choose to STICK at each date, with the variance going to zero. If the initial variance x exceeds x^* the action SWITCH dominates the action STICK. When SWITCH is chosen the variance moves toward x_1^{**} and therefore stays in the interval (x^*, ∞) . Hence the agent chooses the action SWITCH at each date, and the variance converges to x_1^{**} . This is a bifurcation or a catastrophe situation: radically different long-run behavior occurs depending on whether initial variance is less than or greater than x^* . In particular we have the following:

Proposition 3.1. (Overtaking): *Assume Case A. Suppose two agents, I and II, are using the same grade n and have initial posterior variance, x_0^I and x_0^{II} respectively, with $x_0^I < x^* < x_0^{II}$. Then there will be overtaking in the following sense: Agent I (who has lower x) initially has a higher expected output. From some date onwards, II will surpass I in expected output. In particular, the high human capital agent, Agent I, will STICK forever, and the low human capital agent, Agent II, will SWITCH forever.*

3.2. Could case A arise in the presence of competition?

The answer is "Yes", although an agent who always chooses to STICK would eventually be driven out of business by those who chose to always SWITCH. Assume a continuum of farmers that grow corn, each facing the structure in Section 2. Now add a fixed resource for which farmers must compete: To plow his land, a farmer needs a bullock. Bullocks live for one period, and they are supplied

by another sector where there is no technical progress: The technology for raising them is fixed through time, and their price in terms of corn is $P(Q)$, where Q is the aggregate supply of corn. With usual assumptions on preferences, $P(\cdot)$ would be monotonically increasing in Q without bound.⁵ At $t = 0$, suppose the status quo for all farmers is $n = 0$. Farmers are of 2 types, 1 and 2, having measure μ_1 and μ_2 . Suppose that the initial beliefs of type 1's are $x = 0$, while those of type 2's are $x = x_1^{**}$. Then type 1's always STICK, hence their x equals 0 always. Type 2's always SWITCH, hence their x equals x_1^{**} always. If no one were to exit, aggregate output at t would be $Q_t = \mu_1(1 - \sigma_w^2) + \mu_2\gamma^t(1 - x_1^{**} - \sigma_w^2)$. But for a type 1, expected profit, $1 - \sigma_w^2 - P(Q)$, turns negative after some date, and he will optimally plan to exit then. In short, the strategy "always STICK" may be optimal even when others are switching, although it drives the stagnating type 1 out of the market and gives him a long run payoff of zero.⁶

3.3. *Is Case A robust to informational spillovers?*

One should distinguish the purely informational effect of spillovers from their implications for strategic behavior. Concerning the latter, [13], [3] and others have established that the ability to free ride on the results of experiments performed by others leads players to delay experimentation. But for myopic agents the incentive to delay in order to free ride is absent, and so in this case at least, strategic considerations do not matter. On the other hand, the purely informational effect of spillovers acts to change the one-period payoffs. One may well ask: Can any agent ever get stuck if he can see the signals of other agents who are upgrading?

The following example shows that the answer is "Yes". Suppose farmer 1 has used method $n = 0$ for so long that his $x = 0$. One day, farmer 2 moves next door. He knows nothing about $\{\theta_0, \theta_1\}$, and presumes that they are drawn from the steady state distribution implied by (3). Assume first that $\alpha = 0$ (which satisfies condition (i) defining Case A). Then for farmer 2, $\text{Var } \theta_0 = \text{Var } \theta_1 = \sigma_\epsilon^2$. In the first period, 1 will STICK to $n = 0$, and 2 will SWITCH to $n = 1$. Suppose farmers 1 and 2 can see each other's signals. Seeing 2's signal in the next period raises 1's payoff to choosing SWITCH in the following period from $\gamma(1 - \sigma_\epsilon - \sigma_w^2)$ to $\gamma[1 - h(\sigma_\epsilon^2) - \sigma_w^2]$. If this is less than $1 - \sigma_w^2$, 1 will again STICK to $n = 0$. But 2 will again SWITCH, this time to $n = 2$. Because $\alpha = 0$, 2's signals will from this point on be uncorrelated with θ_1 , and so 1's payoff to switching to $n = 1$ remains unchanged. Hence 1 will STICK to $n = 0$ forever. By continuity, this logic remains intact for the case in which α is positive

but small: For informational spillovers to induce the laggard to SWITCH, they must do so early on, before the leader pulls so far away that his signals cease to have any significant effect on the laggard's payoffs.⁷

4. Case B (Positive long-run growth from all initial conditions)

Case B is defined by the reversal of inequalities (i) and (ii) of case A, but with no restrictions on x^* and x_1^{**} :

- i. $\alpha\gamma > 1$ and ii. $G(0,0) < G(0,1)$.

Proposition 4.1 (Case B). Suppose that Case B holds. Then

- (i) (Positive long-run growth) *There exists an $m > 0$ and an integer J such that in every J consecutive periods a jump in grade of size m or higher is chosen at least once. Hence for all dates $t > J$, if k_t denotes the grade at date t under the optimal policy then $k_t/t \geq m/(2J)$.*
- (ii) (The policy functions). *Define $k^*(x)$ to be the set of optimal jumps from x : i.e., $k^*(x)$ is the set of maximizers of $G(x,k)$ over $k \in [0,1]$. Then (a) $\exists \underline{x} \geq 0$ and $\bar{x} > 0$ with $\underline{x} \leq \bar{x} < \infty$ such that $k^*(x) = 1$ for $x \in [0, \underline{x})$ and $k^*(x) = 0 \forall x \geq \bar{x}$ and, (b) either $0 \leq \underline{x} < \bar{x}$ and $k^*(x)$ is single-valued and strictly decreasing in x on (\underline{x}, \bar{x}) or $0 < \underline{x} = \bar{x}$ and $k^*(x) = \{0,1\}$.*

5. A Comparison of Cases A and B

5.1. The role of experience in the two cases

The accumulation of experience with a technology -- i.e., a lowering of x -- in case B promotes upgrading, whereas in Case A it creates a resistance to it, and can cause a subset of agents to experience stagnation in the long run. This difference is highlighted in figure 3, which shows that the policy functions have opposite slopes in each case. The growth literature [7] and [10] emphasizes the growth-enhancing role played by experience, and neglects the possibility Case A raises, i.e., that specialization in a technology can kill long-run growth. The parameters α and σ_ϵ^2 are especially important in determining whether Case A or Case B will arise. If human capital is very technology-specific (small

α and/or big σ_ϵ^2) we will have Case A. But if human capital is fairly general (big α and small σ_ϵ^2), we are likely to have Case B.

Figure 3 around here

5.2. *Do both cases apply in fact ?*

With its implication of occasional upgrading, Case B seems to characterize the bulk of the upgrading behavior of firms -- data on investment in plants show "spikes" at fairly regular intervals, indicating periodic upgrading of technology in these plants [5]. But overtaking behavior is also observed: First, [6] and [9] show that (after controlling for observable characteristics) workers with lower initial wages have higher lifetime wage growth. And second, Korea, Taiwan, Hong Kong, and Singapore have overtaken many slower-growing countries, and their growth rates continue to be above average. Thus Case A also seems relevant.

6. Conclusions

Our paper makes two points. First, human capital accumulation on a given activity, eq. (2b), is linked to how it depreciates when switching to a different activity, eq.(3). The literature [7, eq.s (4.2) and (4.7)], [10, eq.s (2) and (4)], and [16, eq.s (1) and (2)] models them separately, but we show that these two processes are likely to have some common determinants, because α , σ_ϵ^2 , and σ_w^2 enter both equations.

Second, an abundance of knowledge can impede long run progress.^{8,9} This happens in Case A. On the other hand, Case B is like the periodically-upgrading (s, S)-type equilibrium that [10] looks at. In focusing on steady state growth, however, [10], as well as [4] miss the trap of case A. And, unlike the [14] and [1] lock-in, say, in which there is less than a critical mass of human capital, here there is too much.

7. Technical Appendix and Proofs

We define the following

$$(5) \quad \hat{x} \equiv \sigma_\epsilon^2 / (1 - \alpha) \text{ when } \alpha \neq 1 \text{ and } \Psi \equiv 1 - \sigma_w^2 - \hat{x}.$$

For $\alpha \neq 1$, \hat{x} is for each k is the fixed point of $h_2(x, k)$. The term Ψ will be important in the proofs.

$$\begin{aligned} \text{Lemma 7.1.} \quad \partial G(x, k) / \partial k &= \gamma^k \Psi [\ln \gamma] + (\alpha \gamma)^k (\ln \alpha \gamma) (\hat{x} - x) && \text{when } \alpha \neq 1 \text{ and} \\ &= \gamma^k \ln \gamma [1 - \sigma_w^2 - k \sigma_\epsilon^2 - (\sigma_\epsilon^2 / \ln \gamma) - x] && \text{when } \alpha = 1. \end{aligned}$$

Proof: Easy calculus.

Q.E.D.

Proposition 7.2: (i) x_k^{**} increases monotonically in k to x_∞^{**} where $x_\infty^{**} \equiv h_1(\hat{x})$ when $\alpha < 1$ and $x_\infty^{**} \equiv \sigma_w^2$ when $\alpha \geq 1$. (ii) The set $[0, x_\infty^{**}]$ is stable (i.e., regardless of the sequence of jump sizes over time, once the posterior variance enters this set it never leaves it) and is absorbing (i.e., for all $\xi > 0$ and for all initial posterior variances x , the posterior variance will enter the set $[0, x_\infty^{**} + \xi]$ in finite time).

Proof of Proposition 7.2: First suppose that $\alpha < 1$, so $\hat{x} \in (0, \infty)$. One may check that for each $k \geq 0$, $h(x, k)$ is increasing and concave in x . Since for all $k > 0$, $h(\hat{x}, k) \equiv h_1(h_2(\hat{x}, k)) = h_1(\hat{x}) < \hat{x}$ and $h(0, k) > 0$, we conclude that $x_k^{**} \in (0, \hat{x})$. However, for any x in $(0, \hat{x})$, one may check that for all $k' < k''$, $h(x, k') < h(x, k'')$. It is easy to see that this implies that $x_{k'}^{**} < x_{k''}^{**}$. Next, for each k , $h(0, k) \leq x_k^{**} \leq h(\hat{x}, k)$. Part (i) of the proposition then follows by noting that both the right and left hand sides of this inequality converge to $h_1(\hat{x})$ as $k \rightarrow \infty$. Part (ii) follows almost immediately from part (i). Next suppose that $\alpha \geq 1$. Then $h_2(x, k)$ is easily seen to be monotonically increasing in k to $+\infty$ for each x . Since h increases monotonically to σ_w^2 , the proposition then follows for $\alpha \geq 1$. *Q.E.D.*

Lemma 7.3: Suppose $\alpha < 1$ and $\Psi \geq 0$. Then the optimal grade is either $k=0$ or $k=1$.

Proof: From Lemma 7.1,

$$(6) \quad \partial^2 G(x, k) / \partial k^2 = \gamma^k [\ln \gamma]^2 \Psi + (\alpha \gamma)^k (\ln \alpha \gamma)^2 (\hat{x} - x).$$

When $\Psi \geq 0$, $\partial^2 G(x, k) / \partial k^2 \geq 0$ for each x in $[0, \hat{x}]$, so $G(x, k)$ is convex in k for all such x . For each such x , the function $G(x, k)$ therefore attains its maximum over k in $[0, 1]$ at one (or both) of the corners

$k=0$ or $k=1$.

Next, fix an $x \in (\hat{x}, \infty)$. When $\alpha\gamma < 1$, we may use Lemma 7.1 to conclude that $\partial G(x, k) / \partial k > 0$ for all k , so the unique optimal jump size is $k^*(x) = 1$. So suppose $\alpha\gamma > 1$ and $\alpha < 1$. Then one may check that $(1/\gamma^k) \partial G(x, k) / \partial k$ is increasing in k . Hence $\partial G(x, k) / \partial k$ is either everywhere positive, so $k^*(x) = 1$; or $\partial G(x, k) / \partial k$ is initially negative then is positive, in which case the optimal grade is at a corner, 0 or 1; or else $\partial G(x, k) / \partial k$ is everywhere negative in which case $k^*(x) = 0$. Hence the lemma holds *Q.E.D.*

Proof of Proposition 3.1: From the discussion preceding the statement of Proposition 3.1 and Lemma 7.3 it suffices to show that in Case A $\alpha < 1$ and $\Psi \geq 0$. However $\alpha < 1$ follows from the fact that $\alpha\gamma < 1$ and $\gamma > 1$. From Proposition 7.2, $x_1^{**} \leq h_1(\hat{x}) < \hat{x}$. In Case A $x^* < x_1^{**}$, so $x^* < \hat{x}$. In Case A this implies that $G(\hat{x}, 1) - G(\hat{x}, 0)$ is positive. Since $G(\hat{x}, 1) - G(\hat{x}, 0) = (\gamma - 1)\Psi$ the lemma follows. *Q.E.D.*

Lemma 7.4. *Suppose we are in Case B and that*

(7) *when $\alpha < 1$ then $\Psi < 0$.*

For each $x \geq 0$, there exists a $\bar{k}(x) \geq 0$ such that $\partial G(x, k) / \partial k$ is positive, zero or negative according to as k is less than, equal to or greater than $\bar{k}(x)$. Hence $k^(x) = \arg \max_{k \in [0, 1]} G(x, k)$ is well-defined and uniquely defined.*

Proof of Lemma 7.4: Fix any $x \geq 0$ and $k' \geq 0$ and suppose that

(8) $\partial G(x, k') / \partial k \leq 0$.

We will show that this implies that for all $k'' > k'$,

(9) $\partial G(x, k'') / \partial k < 0$.

It should be clear that showing that (8) implies (9) proves Lemma 7.4.

Suppose that $\alpha = 1$. Then from Lemma 7.1, $[\gamma^k \ln \gamma]^{-1} \partial G(x, k) / \partial k = [1 - \sigma_w^2 - (\sigma_\epsilon^2 / \ln \gamma) - x] - k \sigma_\epsilon^2$. Since the right hand side of this inequality is strictly decreasing in k , we see that (8) implies (9) when $\alpha = 1$. Next suppose that $\alpha \neq 1$. From Lemma 7.1, $(1/\gamma^k) \partial G(x, k) / \partial k = [\ln \gamma] \Psi + \alpha^k (\ln \alpha \gamma) (\hat{x} - x)$. By hypothesis, the first term $[\ln \gamma] \Psi$ is negative when $\alpha < 1$. To prove that (8) implies (9) it therefore

suffices to show that

$$(10) \quad \alpha^k (\ln \alpha \gamma) (\hat{x} - x)$$

is either strictly decreasing in k or, when $\alpha < 1$, that it is non-positive. In Case B (where $\alpha \gamma > 1$) (10) is strictly decreasing in k when (i) $\alpha < 1$ and $x < \hat{x}$ or (ii) $\alpha > 1$ (since then $\hat{x} < 0$). Also, (10) is non-positive when (iii) $\alpha < 1$ and $x \geq \hat{x}$. One then checks that (i) - (iii) are exhaustive of the cases where $\alpha \neq 1$ for Case B. *Q.E.D.*

Proof of Proposition 4.1: (ii) When $\alpha < 1$ and $\Psi \geq 0$ we conclude from Lemma 7.3 that $k^*(x) = 1$ for $x \in [0, x^*)$, $k^*(x) = 0$ for $x \in (x^*, \infty)$ and $k^*(x)$ contains both $k = 0$ and $k = 1$ when $x = x^*$. Hence setting $\underline{x} = \bar{x} = x^*$ proves all of part (ii) of the proposition for this situation. So we shall now suppose that whenever $\alpha < 1$, $\Psi < 0$ (i.e., (7) holds). We may use Lemma 7.1 to show that

$$(11) \quad \text{whenever } \alpha \gamma > 1, \quad \partial G(x', k) / \partial k|_{k=0} < \partial G(x'', k) / \partial k|_{k=0} \quad \forall x' > x'' \geq 0.$$

When Case B holds, $G(0, 1) > G(0, 0)$ so $k^*(0) > 0$. Fix any $x' > 0$ and suppose that $\partial G(x', k) / \partial k|_{k=0} > 0$. Then (11) implies that $\partial G(x'', k) / \partial k|_{k=0} > 0$ for all $x'' < x'$. Hence $k^*(x') > 0$ implies that $k^*(x'') > 0$ for all $x'' < x'$. Lemma 7.4 therefore implies the existence of an \bar{x} in $(0, \infty]$ such that part (iia) of this Proposition holds. For x very large it is easy to check that $\partial G(x, k) / \partial k < 0$ for all k when $\alpha \gamma > 1$, so we also conclude that $\bar{x} \in (0, \infty)$. Hence we have an $\bar{x} \in (0, \infty)$ such that $k^*(x) > 0$ for x in $[0, x^*)$ and $k^*(x) = 0$ otherwise.

$G(x^*, 0) = G(x^*, 1)$ so Lemma 7.4 implies that $k^*(x^*) > 0$ so $x^* < \bar{x}$. Eq. (7) implies that when $\alpha < 1$, $\partial G(\hat{x}, k) / \partial k|_{k=0} < 0$ so $\hat{x} > \bar{x}$. For any x such that $k^*(x) > 0$, $k^*(x)$ is the solution to $\partial G(x, k) / \partial k = 0$ unless the boundary $k = 1$ is reached. From Lemma 7.1 it is easy to check that this solution is strictly decreasing in x on $(0, \bar{x})$ whenever $\alpha \gamma > 1$, (where we use the fact that when $\alpha < 1$, $\hat{x} > \bar{x}$, and the fact that when $\alpha > 1$, $\hat{x} < 0$). Hence there will exist an $\underline{x} \geq 0$ such that for x in $[0, \underline{x})$ the boundary $k = 1$ is the optimal solution and for x in (\underline{x}, \bar{x}) , $k^*(x)$ is strictly decreasing in x . This completes the proves of part (ii).

(i) We shall prove this for the situation depicted in Fig. 4 where $0 = \underline{x} < \bar{x}$. The situation where $\underline{x} > 0$ and/or $\underline{x} = \bar{x}$ proved similarly with obvious modifications.

Figure 4 around here

Define the function $D:[0,1] \rightarrow \mathbb{R}_+$ by $D(k)=\bar{x}^{**}$. This function is increasing in k (from Proposition 7.2) and so its inverse function $D^{-1}(x)$ is increasing in x . Also $D^{-1}(0)=0$. From the proof of part (ii) of this proposition, when (7) holds the optimal policy function $k^*(x)$ is decreasing in x with $k^*(0)>0$ and $k^*(x)=0$ for $x>\bar{x}$. Hence (see fig. 4) there exists an intersection point (x_M^{**}, M) for these functions. Fix any m such that $0<m<M$. Define x_m to be such that $k^*(x_m)=m$.

From Proposition 7.2 there exists an $x_\infty<\infty$ such that the posterior variance process enters the set $[0, x_\infty]$ in finite time and stays there forever after. Hence without loss of generality we may suppose that the initial posterior variance lies in this set. From any x in $(x_m, x_\infty]$ a grade less than m is chosen. Since iterates of the function $h(x, m)$ converge to x_m^{**} , and since $x_m^{**}<x_m$, we may choose an integer J such that the J -th iterate $h^J(x_\infty, m) \in [0, x_m]$. From any x in $[0, x_m]$ a grade of m or higher is chosen. Hence the agent will not choose grades less than m for any J or more consecutive periods. Q.E.D.

8. References

- [1] Azariadis, Co. and A. Drazen: "Threshold Externalities in Economic Development," *Quarterly Journal of Economics*, 105 (1991), 501-26.
- [2] Berry, D. and B. Fristedt: *Bandit Problems: Sequential Allocation of Experiments*. London: Chapman and Hall, 1985.
- [3] Chamley, C. and D. Gale: "Informational Revelation and Strategic Delay in a Model of Investment," *Econometrica*, 62, no.5 (September 1994), 1065 - 86
- [4] Chari, V. and H. Hopenhayn: "Vintage Human Capital," *Journal of Political Economy*, 99 no. 6 (December 1991), 1142-1165.
- [5] Cooper, R., J. Haltiwanger, and L. Power: "Machine Replacement and the Business Cycle: Lumps and Bumps." Department of Economics, University of Maryland, July 1994.
- [6] Hause, J: "The Fine Structure of Earnings and the On-the-Job-Training Hypothesis," *Econometrica*, 48 (1980), 1013-1029.
- [7] Lucas, R. E.: "Making a Miracle," *Econometrica* 61, no. 2 (March 1993), 251-72.
- [8] Markoff, J.: "Despite Rivals' Chip Speed, Intel Still Leads PC Market," *NY Times*, vol. CXLIII Sec. D, pp. D1-D2, October 17, 1994.
- [9] Parent, D.: "Matching, Human Capital, and the Covariance Structure of Earnings." Cahier 0995,

Department of Economics, University of Montreal, February 1995.

- [10] Parente, S.: "Technology Adoption, Learning-by-Doing, and Economic Growth," *Journal of Economic Theory*, 63, no.2 (August 1994), 346 - 69.
- [11] Prescott, E.: "The Multi-Period Control Problem Under Uncertainty," *Econometrica*, 40, no. 6 (November 1972), 1043-58.
- [12] Radner, R. and J. Stiglitz: "A Nonconcavity in the Value of Information," in *Bayesian Models in Economic Theory*, ed. by M. Boyer and R. Kihlstrom, New York: Elsevier, 1984.
- [13] Rob, R.: "Learning and Capacity Expansion under Demand Uncertainty," *Review of Economic Studies*, 58, no.4 (July 1991), 655 - 76.
- [14] Romer, P.: "Endogenous Technological Change," *Journal of Political Economy*, 98, no. 5, pt.2 (October 1990), S71 - S102.
- [15] Wilson, R. : "Informational Economies of Scale," *Bell Journal of Economics*, (Spring 1975), vol. 6, No. 1. , 184-95.
- [16] Young, A.: "Invention and Bounded Learning by Doing," *Journal of Political Economy*, 101, no. 3 (June 1993), 443-72.

FOOTNOTES

¹We thank the C.V. Starr Center for Applied Economics at New York University for technical and financial help, and Chung Tse for help with the research. We are grateful for comments from the editor and two anonymous referees.

² Prescott (1972) and Wilson (1975) analyze this type of production function.

³ The information that the agent gets depends on n , but not on z . Hence (2a) and (2b) remain valid even in a multiperiod maximization problem.

⁴ If i. and ii. hold, but if iii fails, we get the uninteresting conclusion that there is no long run growth from any initial condition: After finite time the agent chooses $k=0$ at each date.

⁵ This would occur if, for instance, bullocks could instead be slaughtered for meat, and there was diminishing marginal utility of consuming corn and meat.

⁶ At $t = 0$, a type 1 is better off than a type 2, and there are real-world examples of that. Think of a type 1 as today's Intel corporation. Despite making a slower computer chip (than the rival RISC chip), Intel still holds more than 90% of the chip market for desktop PC's, because popular software is largely available only for Intel-compatible machines (Markoff, 1994). If it doesn't switch to faster chips, however, Intel will eventually lose its market.

⁷ From (3), $y_{n+k} = \alpha^{k/2} \theta_n + \sum_{j=0}^{k-1} \alpha^{j/2} \epsilon_{n+k-j} + w_{n+k}$. If α is small, the signal's informativeness about θ_n decays rapidly as k grows.

⁸ This is not so in the Multi-Arm Bandit model (Berry and Fristedt, 1985) in which there is just one decision: which arm to pull -- which is like choosing n in our model. The essential differences are three: (1) in our model, there is a second decision: how to operate the chosen grade, and (2) learning is not about the grade's quality (which is known), but about how to operate the grade. (3) the Bandit model says that, under discounting, the agent may end up on an inferior arm, but (in contrast to our Case A) it does not imply any tendency for an uninformed agent to overtake an informed one.

⁹ Radner and Stiglitz (1984) show that there may be a noncavity in the value function in a neighborhood of zero information. Our results do not rely on such nonconcavities.

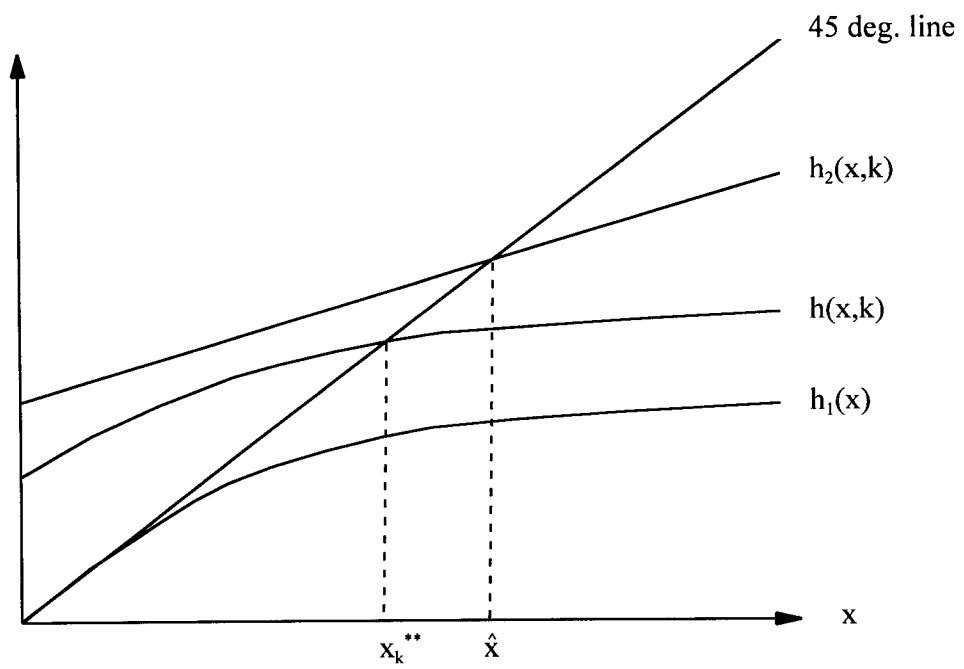


FIGURE 1 - The Functions h_1 , h_2 and h .

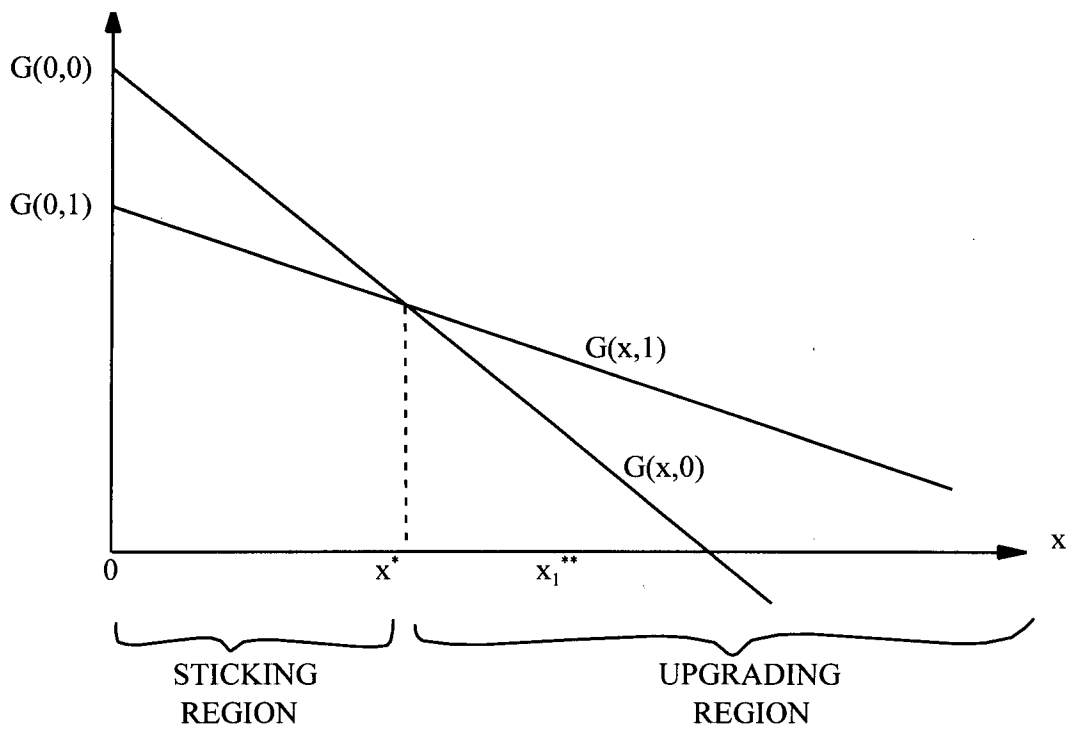


FIGURE 2 - Graph of Case A

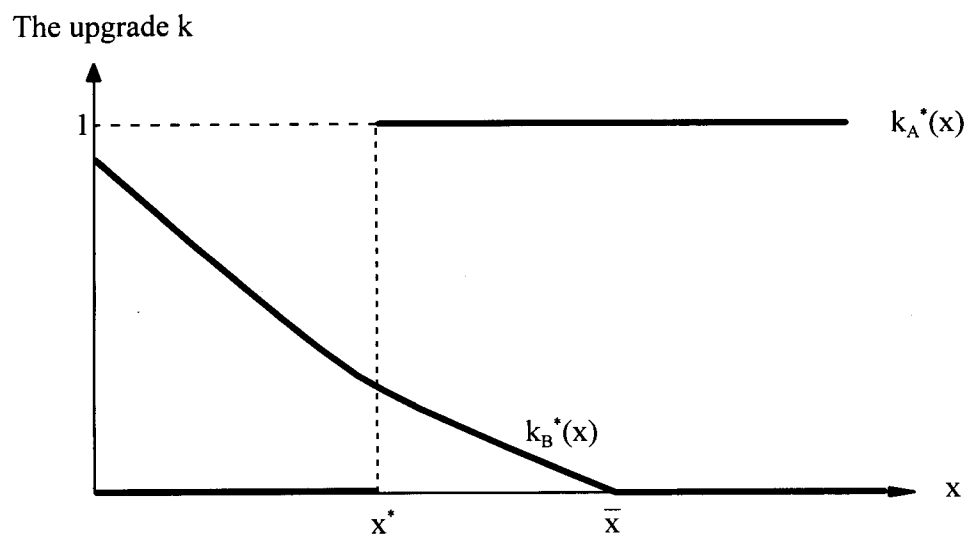


FIGURE 3 - Policy Functions in Cases A and B