NBER WORKING PAPER SERIES

THE BAYESIAN FOUNDATIONS OF LEARNING BY DOING

Boyan Jovanovic
Yaw Nyarko

Working Paper No. 4739

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 1994

We 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. This paper is part of NBER's research program in Productivity. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

NBER Working Paper #4739
May 1994

THE BAYESIAN FOUNDATIONS OF LEARNING BY DOING

ABSTRACT

This paper explores a one-agent Bayesian model of learning by doing and technological choice. To produce output, the agent can choose among various technologies. The beneficial effects of learning by doing are bounded on each technology, and so long-run growth in output can take place only if the agent repeatedly switches to better technologies.

As the agent repeatedly uses a technology, he learns about its unknown parameters, and this accumulated expertise is a form of human capital. But when the agent switches technologies, part of this human capital is lost. It is this loss of human capital that may prevent the agent from moving up the quality ladder of technologies as quickly as he can, since the loss is greater the bigger is the technological leap.

We analyze the global dynamics. We find that a human-capital-rich agent may find it optimal to avoid any switching of technologies, and therefore to experience no long-run growth. On the other hand, a human-capital-poor agent, who because of his lack of skill is not so attached to any particular technology, can find it optimal to switch technologies repeatedly, and therefore enjoy long-run growth in output. Thus the model can give rise to overtaking.

Boyan Jovanovic
Department of Economics
New York University
269 Mercer Street
New York, NY 10003
and NBER

Yaw Nyarko
Department of Economics
New York University
269 Mercer Street
New York, NY 10003

Table of Contents

List of Figures

I. Introduction

This paper explores a one-agent Bayesian model of learning-by-doing and technological choice. In this framework, the more a technology is used, the more productive it gets. Once the productivity gains on a given technology have been achieved, further improvement can be had only by switching to a new technology.

After the switch, how useful an agent's prior experience will be on the new activity will depend on how similar the new activity is to the old. In the Bayesian framework, this depends on how correlated their respective uncertain parameters are. Information that is in this sense transferable fits the notion of general human capital, while information about an unknown parameter that is independent of other unknowns, is exactly like specific human capital.

Our model is like that of Parente (1991), but his formulation is not information theoretic, and he, like Chari and Hopenhayn (1991), looks only at constant growth paths. The outcome that in our model corresponds to steady-state growth is the situation in which the agent switches to a new technology at equally-spaced intervals, and in which the size of the jump (measured in vintages) is the same each time.

An agent can also get stuck in an old technology for ever, and experience no long run growth in output. As in Chari's and Hopenhayn's model, each new technology dominates the previous one at any given level of expertise of its users. The crucial difference is that here, the more remote one's current knowledge is from "frontier knowledge", the less expertise one has with the frontier technology. So, even though the frontier technology keeps getting better, expertise shrinks. Generally, this prevents the agent from upgrading his technology too fast. And if successively better technologies differ enough in the type of expertise needed to operate them properly, an agent will understand them so poorly that he will stick with the old one he understands well. Paradoxically, the type of agent to whom this may happen is one that has learned a particular technology so well that he will not switch to a new, untried one. Such an agent may in the long run be overtaken by an agent that is initially less productive on the technology at hand, and who therefore is more willing to try a new one.

II. Model

A risk-neutral agent can produce a good with one of several technologies indexed by $n = 1, 2, \ldots$. If he uses technology n at date t, a decision z yields net output q via the production function[1]

$$q = \gamma^n [1 - (y_{nt} - z)^2] . \tag{1}$$

Here y_{nt} is a random variable that acts as an unknown "target", and is observed after z is chosen. Since $\gamma \geq 1$, a larger n denotes a better technology. Let $E_t(\cdot)$ denote the conditional expectation at date t. The decision that maximizes $E_t(q)$ in (1) is

$$z = E_t(y_{nt}) . \tag{2}$$

The random target fluctuates around a technology-specific parameter θ_n:

$$y_{nt} = \theta_n + w_{nt} . \tag{3}$$

The agent does not know θ_n. He can observe y_{nt}, but only if at date t he uses technology n.[2] Assume that w_{nt} is an i.i.d variate, with mean zero and variance σ_w^2. Since $E_t(w_{nt}) = 0$, equation (2) implies that the optimal decision is

---

[1] This type of production function has been analyzed by Prescott (1972) and Wilson (1975).

[2] The information that the agent gets depends on what vintage he chooses to operate, but not on the value of z. Therefore equation (2) holds even in a multi-period maximization problem.

$$z = E_t(\theta_n) , \tag{4}$$

and (1), (3) and (4) imply that expected net output is

$$E_t(q) = \gamma^n [1 - \mathrm{Var}_t(\theta_n) - \sigma_w^2] , \tag{5}$$

where $\mathrm{Var}_t(\cdot)$ denotes the conditional variance. If he uses technology n, he also observes y_{nt} and learns more about θ_n, which allows him to make a better decision z. This reduces $\mathrm{Var}_t(\theta_n)$, and raises his expected net output. However this learning process is <u>bounded</u>: Using technology n 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.

III. The Transfer of Human Capital.

There is no direct cost of switching to a different technology, and no adjustment costs in z. The only link between technologies is informational. Let

$$\theta_{n+1} = \sqrt{\alpha}\theta_n + \epsilon_{n+1} \tag{6}$$

for all n, where ϵ_n is an i.i.d. normal variate with mean zero and variance σ_ϵ^2. If the agent has not yet tried technologies $n+1$, $n+2$, ..., equation (6) implies that

$$\mathrm{Var}_t(\theta_{n+1}) = \alpha \mathrm{Var}_t(\theta_n) + \sigma_\epsilon^2 . \tag{7}$$

The earlier equations in the model allow us to think of the posterior precision on θ_n as an index of human capital. Equation (6) adds a further dimension and lets us evaluate some hypotheses that are connected to human capital -- general and specific. If $\alpha = 1$ and $\sigma_\epsilon^2 = 0$, the correlation coefficient between θ_n and θ_k is unity, for all k, which corresponds to the case where human capital is general and freely transferable across technologies. But if $\alpha = 0$, human capital is technology-specific.

We choose the AR-1 formulation in (6) to model the transfer of knowledge because by varying α, we capture many types of evolution of technologies:

i. When $0<\alpha<1$, note that the θ_n process has a tendency toward zero;

ii. When $\alpha=1$ then $E\theta_{n+1}=\theta_n$, so the technology parameter tends to where it was in the previous vintage.

iii. When $\alpha>1$ then there is drift towards infinity.

Notice that α has the interpretation of a human capital transfer parameter. For let σ_n^2 denote the variance of the parameter θ_n. Then from (6), $\sigma_{n+1}^2=\alpha\sigma_n^2+\sigma_\epsilon^2$. We may therefore write $\Delta\sigma_{n+1}^2=\alpha\Delta\sigma_n^2$ where $\Delta\sigma_n^2$ represents a "change in σ_n^2." Hence a unit change in human capital on technology n results in α units of change in the human capital on technology n+1.

We assume that the prior over θ_1 at date 1 is normally distributed. Eqs. (3) and (6), and the normality assumptions made on w_{nt} imply that the posterior belief at each date over the parameter of any vintage, θ_n, will also be normally distributed.

We define the following functions of x:

$h_1(x) \equiv \sigma_w^2 x/(\sigma_w^2 + x)$; *(updating)*

$h_2(x) \equiv \alpha x+\sigma_\epsilon^2$; and *(transfer of knowledge)*

$h(x) \equiv h_1(h_2(x))$. *(transfer and updating)*

These functions have the following interpretations: Suppose that vintage n is the newest vintage that the agent has worked with at the <u>end</u> of date t and suppose the agent has a posterior of $x_{n,t}$ over θ_n. Then $h_2(x_{n,t})$ is the agent's <u>prior</u> at date t over θ_{n+1}. Suppose the agent uses vintage n at date t+1. The agent will then observe $y_{n,t+1}$ after which the posterior variance over vintage n becomes, via Bayesian updating, $h_1(x_{n,t})$. Suppose instead that at date t+1 the agent chooses vintage n+1. The agent will then observe $y_{n+1,t+1}$. The agent's posterior variance over θ_{n+1} will then be $h(x_{n,t})$.

We define $\hat{x}$ to be the fixed point of the h_2 map. This fixed point exists and is unique whenever $\alpha < 1$; when $\alpha \geq 1$ we take $\hat{x}$ to be $+\infty$. We define x** to be the fixed point of the h map. Since h is as drawn in Figure 1 (i.e., since $h(0) > 0$ and h is bounded and concave), x** exists and is unique.

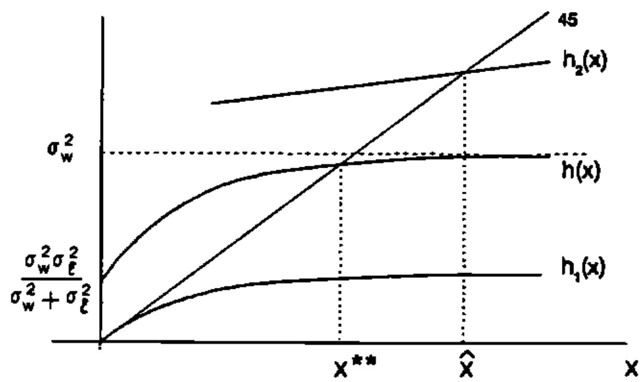

Figure 1: <u>The Mappings h_1, h_2 and h</u>.

The following property of h will be used later: Suppose that at the end of date t the agent has a posterior over a "status quo" vintage n equal to x_0. Suppose further that at each date he switches to the next technology. Then his posterior over the most recent vintage is given by iterates of h from x_0. That is, his posterior over θ_{n+1} at the end of date t+1 is $x_1 = h(x_0)$, his posterior over θ_{n+k} at the end of date t+k is the k-th iterate of h, $x_k = h(h(\ldots(h(x_0))$ (k-times), and so on. For any x_0, the sequence $\{x_k\}^{\infty}_{k=1}$ converges <u>monotonically</u> to x**.

a. The Human Capital Transition Equation. Several authors have discussed the question of how human capital at date t depends on the nature of prior experience, and on the content of previous human capital. The specific functional form implied by our model depends in a crucial way on how we define things.

One way to think of this equation is in linking per unit productivity at date $t+1$ when vintage $n+1$ is being used for the first time, with the history of outputs from dates 1 through t. This is, formally speaking, the approach in Lucas (1993). Since in our model there is a one-to-one relationship between productivity and variance this is equivalent to looking at the way in which the variances at those dates relate to each other. In our model the relationship is Markov. The date $t+1$ variance of vintage $n+1$ is a function of only the variance at date t of vintage n. In particular the date t variance of vintage $n+1$ is related to the history of the variances of all previous vintages by the Markov relationship $x_{t,n+1}=h_2(x_{t,n})$.

An alternative way of thinking about the human capital transition equation is look for the function linking the <u>experience</u> in previous vintages to the productivity on the current vintage. This is the language Lucas actually uses. In particular suppose that at some date T, vintage 1 was used for the first τ_1 periods, then vintage 2 was used for the next τ_2 periods, ..., and vintage n-1 was used for τ_{n-1} periods. Then the human capital equation will have the productivity on vintage n equal to a function of the numbers τ_1 through τ_{n-1}:

$$\text{Variance on Vintage } n = x_n = h_2(h_1^{\tau_{n-1}}\ (h_2(h_1^{\tau_{n-2}}\ (\ldots\ h_2(h_1^{\tau_1}\ (x_1))))), \tag{8}$$

where x_1 is the initial date 1 prior variance over the date 1 status quo technology, and where, via Bayes' rule

$$h_1^{\tau}(x) = \sigma_w^2 x/[\sigma_w^2 + \tau x] \tag{9}$$

gives the effect of τ units of experience on the variance of a vintage.

We now show that a geometrically-declining weights model can be obtained as an approximation of our model. Although in this exercise τ is an integer, h_1^{τ} is defined for all τ in $[0,\infty)$. Now differentiate (8) with respect to τ_k where h_1^{τ} is as in (9). This differentiation exercise asks: if we raise by a unit the experience level on vintage k, what is the effect on the human capital level (or more appropriately, the variance of) vintage n?

$$\partial x_n/\partial \tau_k = \alpha^{n-k}.h_1^{\tau_{n-1}}(x_{n-1}).h_1^{\tau_{n-2}}(x_{n-2})\ldots h_1^{\tau_1}(x_k), \tag{10}$$

where x_r is the argument of $h_1^{\tau_r}$ in (10) above and denotes the variance of vintage r when vintage r is just about to be used for the first time. Suppose that to a first approximation,

$$\tau_k = \tau_{k+1} = \ldots = \tau_n \text{ and } x_k = x_{k+1} = \ldots = x_n (= x' \text{ say}).$$

Then from (9), $h_1^{\tau_r}(x_r)$ is equal to a constant, ξ say, for all r. So from (10) we obtain $\partial x_n/\partial \tau_k = (\alpha\xi)^{n-k}$. This in turn implies that, to a first approximation,

$$x_n = \sum^{n-1}_{r=1} (\alpha\xi)^{n-r}\tau_r \ .$$

Hence the productivity measure for vintage n is a function of the experience levels of agents with geometrically declining weights. This is of a form similar to that of eq. (4.7) of Lucas (1993). We obtain the same general formulation as in Lucas if $(\alpha\xi)<1$. Here, however, $\alpha\xi$ can exceed unity, in which case vintages further in the past have a stronger effect on today's output than more recent vintages. This of course can happen only when $\alpha>1$.

b. The Parente and LPSY Conditions. Parente (1991) assumes that when an agent switches to <u>any</u> other activity, his human capital depreciates, and the bigger is the jump the greater is the fall in human capital. In our framework if x is the posterior variance on the current technology (vintage n, say) then $h_2^k(x)$ (the k-th iterate of the function h_2) is the prior variance over vintage n+k. The depreciation of human capital is therefore given by

$$D(k) = h_2^k(x)-x.$$

From figure 1 whenever $x<\hat{x}$, $D(k)>0$ for all $k=1,2,\ldots$ Further, $D(k)$ is bigger the bigger is k. Hence the Parente assumption is equivalent to the following:

Defn: The Parente condition, or condition P, holds at x if $x<\hat{x}$.

If we define $x_{Parente}=\hat{x}$, then condition P holds at each $x<x_{Parente}$.

Lucas (1993), Parente (1991), Stokey (1991) and Young (1993) emphasize a different condition - that human capital in some future vintage will be larger if today a higher rather than a lower vintage is chosen. For example, suppose that $N>n_2>n_1$ and tomorrow an agent will be using vintage N. Then human capital in vintage N tomorrow will be higher if vintage n_2 is used today than if vintage n_1 is used. Surprisingly, this is not always true in our model. Fix an n and set $n_1=n$ and $n_2=n+1$. Consider two agents with the *same* knowledge in the sense that each has $Var_t(\theta_n) = x$. We shall call them agent S and agent F (for "Slow" and "Fast" resp.). Assume that in period $t+1$ agent S operates technology n while agent F operates technology $n+1$. Which agent will be better prepared for technology $n+2$ in period $t+2$? Well, let $\phi_S(x)$ (resp. $\phi_F(x)$) be the variance of agent S's (resp. agent F's) beliefs over θ_{n+2}:

$$\phi_S(x) \equiv h_2(h_1(x)) \quad \text{and} \quad \phi_F(x) \equiv h_1(h_2(x)).$$

Defn: The Lucas-Parente-Stokey-Young, or condition LPSY, holds at x if $\phi_S(x) > \phi_F(x)$. We also define $x_{LPSY} = \text{Sup}\ \{x \geq 0 \mid \text{ the LPSY condition holds at } x\}$.

When LPSY holds, agent F is better prepared for technology $n+2$ from x. The shapes of the ϕ_S and ϕ_F functions and x_{LPSY} are illustrated in figures 2 and 3 below and described in the lemma below. Surprisingly, we see that it is indeed possible for condition LPSY to fail.

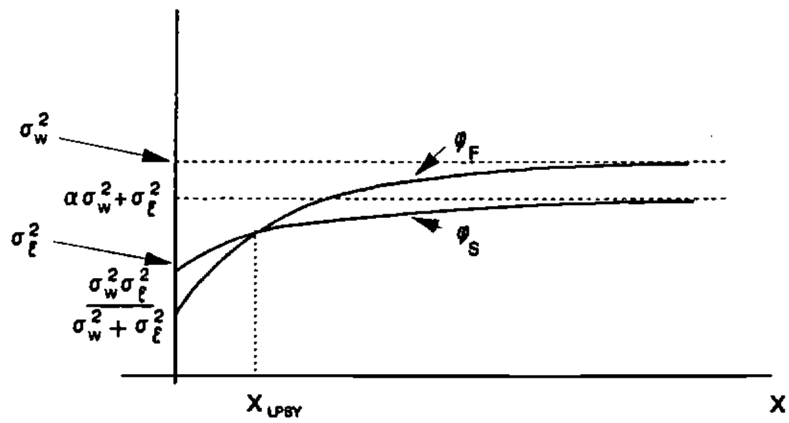

Figure 2: LPSY Fails for Large x.

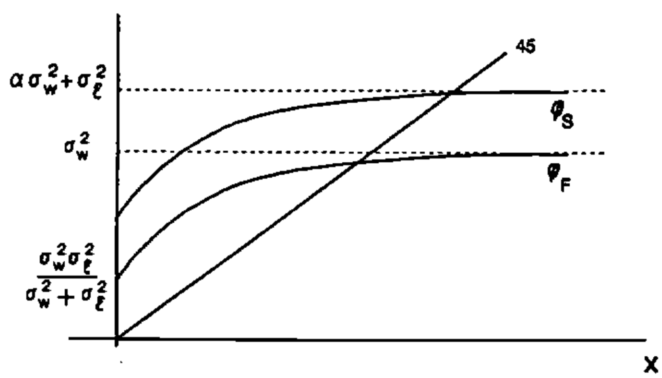

Figure 3: LPSY Holds for all x.

Lemma 3.1.

(i) When $\hat{x} < \sigma_w^2$, $0 < x_{LPSY} < \infty$, the LPSY condition holds at each $x < x_{LPSY}$ and is violated at each $x \geq x_{LPSY}$.

(ii) When $\hat{x} \geq \sigma_w^2$, and hence in particular when $\alpha \geq 1$, $x_{LPSY} = \infty$ and the LPSY condition holds at each $x \geq 0$.

(iii) $\hat{x} < x_{LPSY}$. In particular the LPSY condition holds at each $x \leq \hat{x}$.

We may refer to our earlier condition LPSY as a "one-step ahead LPSY," since it involved comparisons of the use of a vintage n with a vintage n+1. Below is the "many-step" ahead version - a kind of first-order dominance result in vintage levels - which is implied by condition LPSY:

Proposition 3.2. Suppose that at date 1 agents S and F both start from the same prior variance over the initial date 1 status quo technology, and this obeys the LPSY condition. Suppose that at each of dates $t=1,2,\ldots,T$, agent S chooses a technology of lower (or equal) vintage than agent F, and suppose that their choice differs at least in one period. Let N be any technology of vintage greater than or equal to the maximum that any of the two agents have chosen by date T. Then at date T+1 agent F will have more human capital (i.e., lower posterior variance) on vintage N than does agent S.

From any given x, neither condition P nor LPSY necessarily holds. (See fig. 2). However, part (iii) of lemma 3.1 means that Condition P implies Condition LPSY. (Equivalently, $x_{Parente} \leq x_{LPSY}$.) The reverse claim, namely that condition LPSY implies condition P, is false -- as figure 2 shows.

In summary we see that if $\alpha < 1$ and x is large both the Parente and the LPSY condition fail. If $\alpha \geq 1$ then both $x_{Parente}$ and x_{LPSY} equal ∞ so the two conditions hold at each finite x. Let us now pursue the Parente condition a little further. When $x > \hat{x}$ (which can only happen when $\alpha < 1$) iterates of the h_2 map from any x are decreasing. This means that the prior variances over later vintages are decreasing the further out is the vintage. The agent has higher human capital on vintages which are further away! This of course is precisely the meaning of the violation of the Parente condition.

One may be dissatisfied with our framework because it allows for this potentially non-intuitive feature. However, in that case, it should be stressed that what one is really dissatisfied with is the assumption that $\alpha < 1$ and x is large. For note what this implies: When x is large, the variance of ϵ, σ_ϵ^2, is relatively small so for each n, $\theta_{n+1} = \alpha^{1/2}\theta_n$ approximately. A large x implies that θ_n is large relative to its mean. Hence when $\alpha < 1$, θ_{n+1} is a fraction of θ_n so with high probability will be much smaller than θ_n so will have a smaller variance. The assumption that $\alpha < 1$ implies that the process $\theta_{n+1} = \alpha^{1/2}\theta_n + \sigma_\epsilon^2$ has a tendency toward zero when the noise term ϵ is small. This explains the non-intuitive feature mentioned earlier. When $\alpha \geq 1$ the process no longer tends towards zero and the non-intuitive feature (i.e., the violation of the Parente condition) disappears. Further, although condition P does not hold for $x > \hat{x}$, the set of beliefs $[0, \hat{x}]$ is in fact absorbing -- starting from any initial belief in this set and following any policy implies that beliefs always remain in this set. Moreover, since x** is strictly less than $\hat{x}$, it is also easily shown that under any policy, starting instead from any initial beliefs exceeding $\hat{x}$, beliefs will enter the set $[0,\hat{x}]$ in finite time, and remain in it thereafter. (This can be seen by observing the shapes of the h_1, h_2 and h maps and noting that the beliefs of any vintage at any date are iterates of some combinations of these maps.)

c. **Switching Constraints.** We now study the dynamics of our model. There are at least three interesting ways of constraining technological switches:

i. *(No Jump Model)* The first is to assume that the agent can not skip intermediate vintages when switching, so that if he wishes to advance from vintage n, he must use vintage n+1 before he can use any higher vintage. We shall refer to this as the "No Jump Case". This case is the easiest to analyze, but it begs the question of why it is impossible to skip vintages when switching.

ii. *(Chari-Hopenhayn Model)* The second type of constraint on switching allows jumps, but only up to some frontier that advances exogenously. This is the case that Chari and Hopenhayn analyzed. This switching constraint is only partially satisfactory in that the exogenous outward movement of the frontier, similar to the Solow growth model, is left unexplained.

iii. *(Full-Menu Model)* The third approach is to leave the choice of vintage unconstrained. This is the approach that Parente took, and we refer to it as the "Full Menu Case". Although the rate of growth is endogenous in all three cases, only this third case is free of any arbitrarily imposed constraint on the rate at which future technologies are chosen.

Section IV will cover the "No Jump Case", section V will take up the "Full Menu Case", and section VI reports results that relate the two cases. Both cases will be analyzed under the following constraint:

The no recall constraint: Once a vintage has been passed over for a higher vintage, it is never recalled. Hence an agent who chooses a vintage n at date t, can not choose a lower vintage $n' < n$ at any future date $t' > t$.

We prove in the appendix that, at least for the myopic and the two-horizon versions of the model, it is actually not optimal to ever exercise such a recall option.

IV. Myopically Optimal Time Paths in the "No Jump" Case.

This section analyzes the optimal behavior of a myopic agent who can upgrade his vintages only one at a time - the No Jump model. Our reason for starting the analysis with myopically optimal policies and time-paths is that first, they are the simplest to characterize, and that second, they share similar broad features with dynamically optimal policies and time-paths.

At any date the agent has a status quo technology - the most recent one that he tried. Define the status quo technology for the initial period, date 1, to be vintage 1. At each date there also is a frontier technology, one vintage higher than the status quo vintage. That is, if the status quo technology is vintage n, the frontier technology is vintage n+1. We refer to the decision to use the status quo technology as "NO SWITCH", and the decision to move to the frontier technology as "SWITCH". At each date the agent must choose between the status quo vintage and the frontier vintage. He can not recall vintages earlier than the status quo vintage. "The variance of vintage n" will mean the variance of θ_n. At the beginning of date 1 the agent has a posterior variance x over vintage 1. This is his "prior belief" about vintage one.

Define

$$\ell(x) \equiv 1 - x - \sigma_w^2\,, \quad \text{and} \quad r(x) \equiv \gamma[1 - \alpha x - \sigma_\epsilon^2 - \sigma_w^2].$$

Once chosen, the optimal "organization" of a vintage is given in (4). Therefore when he has a prior belief x over vintage n, $\gamma^n\ell(x)$ is the agent's expected net output on vintage n, and $\gamma^n r(x)$ is his expected net output on vintage $n+1$. Let $x_{n,t}$ denote his posterior variance over θ_n at date t. Let x^* be the unique point where $\ell(x)$ crosses $r(x)$; i.e., $\ell(x^*) = r(x^*)$.

The immediate gain to switching technologies is

$$\Delta = r(x) - \ell(x) = \gamma - 1 + (1-\alpha\gamma)x - \gamma\sigma_\epsilon^2 - (\gamma-1)\sigma_w^2.$$

There are two parameters -- α and σ_ϵ^2 -- that positively affect the depreciation of human capital by raising the variance of θ on the frontier technology, and each parameter indeed does reduce Δ. The third parameter that reduces Δ is σ_w^2. One certainly would have expected an increase in this parameter to slow down the rate at which the frontier technology is learned, and therefore to reduce the dynamic incentive to switch, but it is somewhat surprising that this parameter

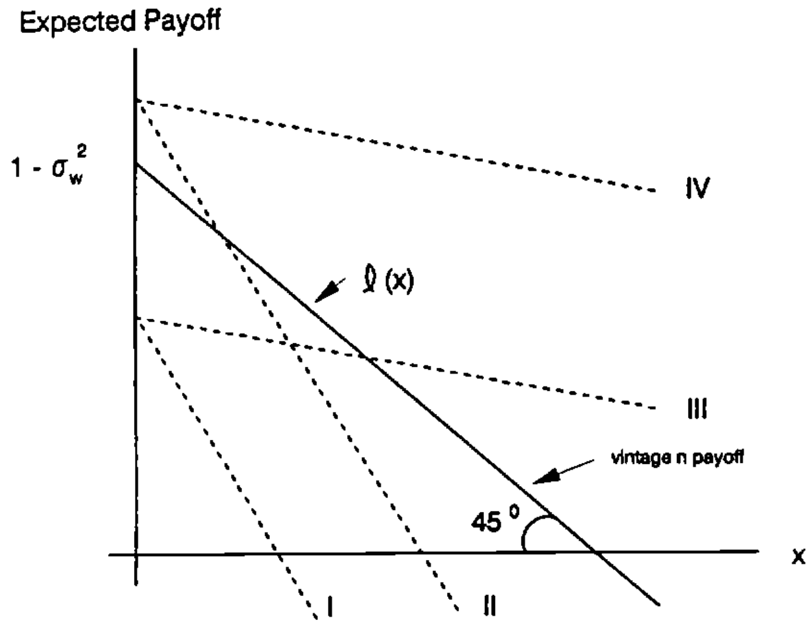

Figure 4: The Payoffs to Using the Current or the New Technologies.

	$\gamma\alpha > 1$	$\gamma\alpha < 1$
$1 - \sigma_w^2 > \gamma(1 - \sigma_\epsilon^2 - \sigma_w^2)$	I Choose Old Vintage Always	III Initial Beliefs Affect Long-Run Growth
$1 - \sigma_w^2 < \gamma(1 - \sigma_\epsilon^2 - \sigma_w^2)$	II Cycles or eventually always switch	IV Switch to Next Vintage at Each Date

Figure 5: The Four Possible Regimes.

should reduce the immediate returns too. The reason why it does so in our model is that the amount by which a rise in σ_w^2 reduces the expected output of a technology is greater for the frontier technology than for the status quo technology. So the immediate and the dynamic effects of a rise in σ_w^2 both work in the same direction.

When γ increases, the effect on Δ is ambiguous. This is because a higher γ raises expected net output at lower levels of x, but it actually lowers expected net output at high levels of x. An increase in γ therefore raises the incentive to switch at low levels of x and lowers it at high levels of x.

Finally, the effect of a rise in x on Δ is also ambiguous, and depends entirely on whether $\alpha\gamma$ bigger or less than unity. This condition determines whether high human capital (i.e., a low x) makes switching more attractive, and the answer is that it does so only if $\alpha\gamma$ exceeds one. When this condition is met, an extra unit of human capital raises output on the frontier technology by more than it raises the output on the status quo technology. Whether this condition is likely to be met in a given technological area will depend on how technologically linked the successive vintages are -- this determines the magnitude of α. For example, if each generation of computer chips is defined as a different technological vintage, then because each generation of chip builds on the previous one, α should be high.

With these remarks out of the way, we can now analyze myopically optimal paths. There are four regimes; the parameters determined which one will prevail. These regimes are determine by whether $\alpha\gamma$ is less than or greater than one, and whether r(0) is greater than or less than $\ell(0)$. In Figure 4, the solid line is $\ell(x)$, while the dashed lines depict four possible schedules for r(x). (In all of this discussion we ignore the non-generic cases where $\alpha\gamma=1$ and/or $\ell(0)=r(0)$.)

Analysis of the Four Cases for the No Jump Myopic Model:

Case I (Agent is stuck at old vintage): In this case, $\ell(x) > r(x)$ for all $x \geq 0$, so that whatever the value of $x_{n,t}$ at the end of date t, vintage n yields a higher expected profit than vintage n+1. He therefore chooses vintage n which causes $x_{n,t}$ to decrease. But, since $\ell(x)$

> $r(x)$ for all x, vintage n will remain preferred over vintage $n+1$ at date $t+1$, and indeed in each subsequent period. Hence the agent will choose the older vintage n at each date, and $x_{n,t}$ converges to zero.

Case IV (Agent always "switches"): In this case, $r(x) > \ell(x)$ for all $x \geq 0$, so that whatever the posterior variance, x, over the parameter θ_n of any vintage n at the end of any date t, vintage $n+1$ yields a higher expected net output than vintage n. Hence he will always switch to the newest vintage at each date, and x converges to x^{**}. Log output will then have a trend on $\ln \gamma$, and its deviations of around this trend will be i.i.d.

Case II. In this case, $x < x^*$ implies $\ell(x) < r(x)$, and $x > x^*$ implies $\ell(x) > r(x)$. Suppose that at date t, the agent must choose between a status quo vintage n and a frontier vintage $n+1$, and let $x_{n,t}$ be his posterior over θ_n.

First, let $x_{n,t} > x^*$. Then he chooses vintage n. He then sees $y_{n,t}$, which lowers his posterior variance over θ_n. He optimally uses vintage n until the variance over vintage n falls below x^*, which it eventually must. Suppose that this occurs after τ periods, at date $T = t+\tau$. Since the posterior over θ_n is then less than x^*, he chooses vintage $n+1$ at that date.

Second, let $x_{n,t} < x^*$. Then he chooses vintage $n+1$ immediately at t. This is equivalent to the setup of the previous paragraph when $\tau = 0$ so that $T = t$. So from now on we may suppose that we are at a date T with $x_{n,T} < x^*$, with $T = t+\tau$, and where τ is an integer and may be zero.

To determine what happens in the subsequent periods, i.e., from date $T+1$ on, we consider two sub-cases: $x^* < x^{**}$, and $x^* > x^{**}$.

Case IIA: $x^* > x^{**}$ (Use inferior technology for $\tau \geq 0$ periods, then "switch" forever.)

Assume $x^* > x^{**}$. We saw earlier that at some date T the agent will have a posterior, $x_{n,T}$, over θ_n smaller than x^*. He then chooses vintage $n+1$, sees y_{n+1}, and his posterior over θ_{n+1} becomes $x_{n+1,T+1} = h(x_{n,T})$. Since iterates of h converge monotonically to the x^{**}, $x_{n+1,T+1}$ is closer to x^{**} than $x_{n,T}$ was. Since both x^{**} and $x_{n,T}$ are less than x^*, so is

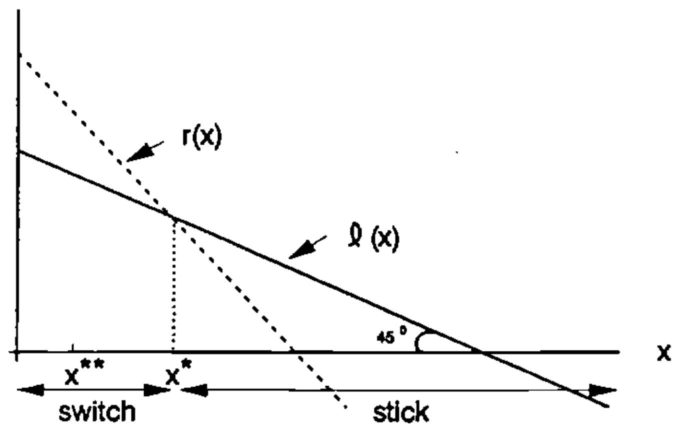

Graph of Case IIA

$x_{n+1,T+1}$. Hence, at the end of T+1 he switches to the frontier technology, n+2. Repeating this logic we see that at each subsequent date he switches to the newest frontier technology.

To summarize Case IIA: There is an integer τ (which could be zero) such that the agent chooses the inferior vintage n for τ periods. Thereafter he switches to the frontier technology at each date.

Case IIB: $x^* < x^{**}$ ("Cycles").

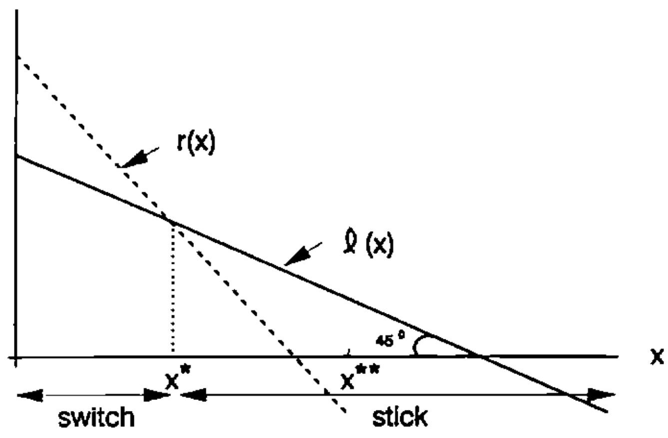

Graph of Case IIB

Again, at date T the agent's posterior over θ_n is $x_{n,T} < x^*$. He then switches to vintage $n+1$, and the posterior over θ_{n+1} is $x_{n+1,T+1} = h(x_{n,T})$. Since iterates of h converge monotonically to x^{**}, this means that $x_{n+1,T+1} > x_{n,T}$. If $x_{n+1,T+1}$ is still less than x^*, he switches to vintage $n+2$, and $x_{n+2,T+2} = h(x_{n+1,T+1})$ is even closer to x^{**}. He will keep switching until the posterior gets larger than x^* which it eventually must since iterates of h converge to x^{**}.

He then begins to forsake frontier technologies and stays with the relatively inferior vintage. But then we are back where we started in CASE II: he stays with the inferior technology until its variance falls below x^*. He then switches to the frontier technology. He then keeps switching to each latest frontier technology, the variance of which continually increases towards x^{**} and eventually exceeds x^*. So we get "cycling": The agent "sticks" with an inferior technology for a while, then continually "switches" to the newest frontier one, then "stays" with an inferior one for a while, then continually "switches" to the newest one, etc, etc, ad infinitum.

An example with parameter values that satisfy the restrictions of case II is described in figure 6. In this example, $x^* = 0.129$, and therefore for any smaller value of x it is optimal to SWITCH. In particular, this is optimal at the value $x_0 = .0896$. Having used the new technology, say technology $n+1$ for one period, it remains optimal for the agent to use it for one additional period, because $h(x_0)$ exceeds x^*, as shown in figure 6. But after two periods of using technology $n+1$, the agent's posterior variance on $n+1$ is once again x_0, and so it is now optimal for him to switch to technology $n+2$, where the whole process repeats itself, and so on[3].

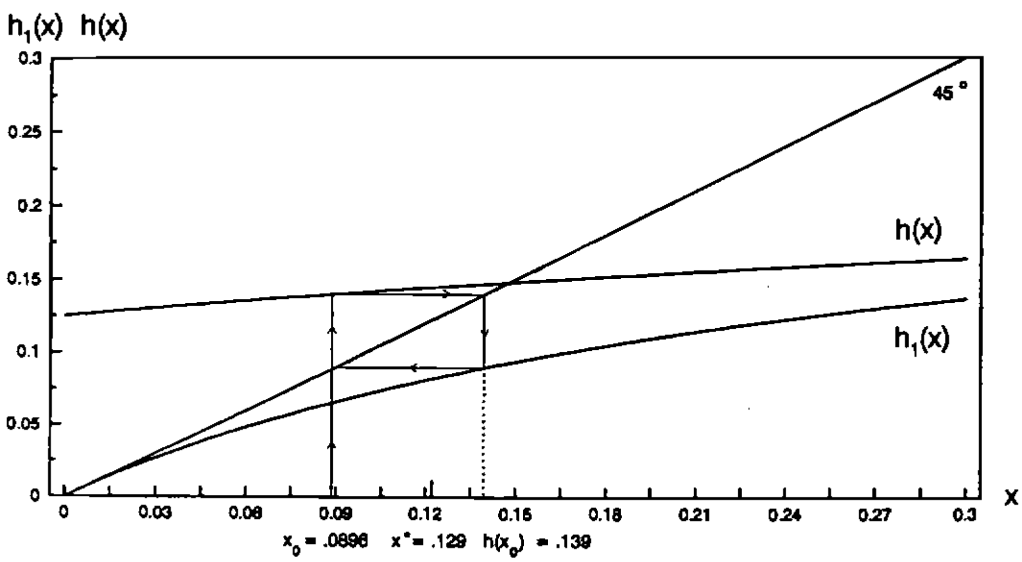

Figure 6: A Two-Period Cycle.

[3] The parameter values underlying this example were $\sigma_w^2 = \sigma_\epsilon^2 = 0.25$, $\gamma = 1.54$ and $\alpha = 0.75$.

<u>Case III</u>. In this case, $x > x^*$ implies $\ell(x) < r(x)$ and $x < x^*$ implies $\ell(x) > r(x)$. Suppose that at t, the frontier vintage is $n+1$, and the "status quo" vintage is n. Let $x_{n,t}$ again be the posterior over θ_n.

First, let $x_{n,t} < x^*$. Then the agent chooses vintage n, sees y_n, and his variance on θ_n falls and therefore remains below x^*. Hence he uses vintage n for ever even though vintage $n+1$ is always available.

Second, let $x_{n,t} > x^*$. Then he uses the frontier vintage $n+1$, and then sees y_{n+1}. Let $x_{n+1,t+1}$ be the posterior over θ_{n+1} at the end of date $t+1$. What happens then depends on whether $x_{n+1,t+1}$ remains larger than x^* or not. This in turn depends on the map of h. We need to consider two sub-cases:

<u>Case IIIA: $x^* < x^{**}$ (Beliefs affect long run growth catastrophically)</u>.

Suppose $x_{n,t} > x^*$ so that at t the agent switches to the frontier technology. Now $x_{n+1,t+1} = h(x_{n,t})$ is closer to x^{**} than $x_{n,t}$ was. Since $x_{n,t}$ and x^{**} exceed x^*, so does $x_{n+1,t+1} = h(x_{n,t})$. Hence at the end of date $t+1$, after using vintage $n+1$, he switches to vintage $n+2$. Repeating this logic we see that if $x_{n,t} > x^*$, he will at each date switch to the newest frontier technology.

To summarize Case IIIA: If initial beliefs are less than x^*, the agent chooses the inferior technology at each date. But if they exceed x^*, he will at each date switch to the newest frontier technology. This is a bifurcation or a catastrophe situation: radically different long-run behavior occurs depending on whether initial posterior variance is less than or greater than x^*.

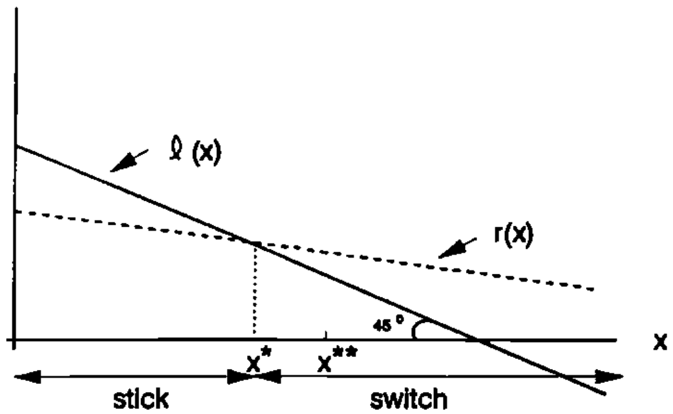

Graph of Case IIIA

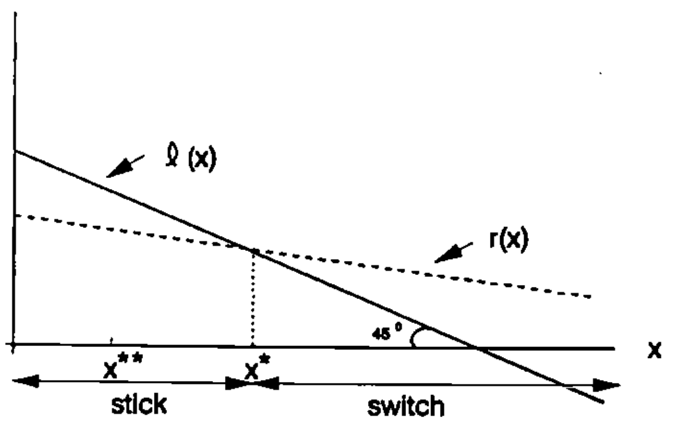

Graph of Case IIIB.

Case IIIB: $x^* > x^{**}$ (Switch for $\tau \geq 0$ periods then "stick" with a vintage).

Suppose that $x_{n,t} > x^*$. Since $\ell(x_{n,t}) < r(x_{n,t})$, the agent will then choose vintage $n+1$ at date $t+1$. The posterior variance of the agent is given by h, whose iterates from $x_{n,t}$ decrease monotonically to x^{**}. Hence he switches to the frontier vintages for a finite number of periods $\tau \geq 1$, after which the posterior variance falls below x^*. The date will then be $T = t+\tau$ and the vintage $N = n+\tau$ will have just been used. But when $x < x^*$, $\ell(x) > r(x)$, so he will prefer vintage $N = n+\tau$ to vintage $N+1$ at date $T+1$. The posterior variance on vintage N then decreases, and so remains below x^*. Hence he will use vintage N at $T+2$ as well. Repeating this logic we see that he will use vintage N in each subsequent period.

To summarize case IIIB: Depending upon the initial belief, the agent will "switch" to a new frontier technology in each of a finite number of periods $\tau \geq 0$. If the initial belief is in the set $[0, x^*)$, then $\tau = 0$. Otherwise $\tau > 0$. After the τ periods, he will remain with the status quo technology at that date (vintage $n+\tau$) forever thereafter. Hence there is some initial "growth" in the vintages, after which there is "no growth" in the vintages.

Summary of all the Cases.

We now list the types of behavior that arise, roughly in increasing order of long-run growth rates. At the extremes, staying with the oldest vintage is the "slowest growth rate" while always switching is the "highest."

Case I	"Stick" to an old vintage forever.
Case IIIB	"Switch" for $\tau \geq 0$ periods then "stick" with a vintage.
Case IIIA	Beliefs affect long run growth. When initial posterior variance is low agent is stuck with initial old vintage; when initial posterior is high, he switches to the frontier technology at each date.
Case IIB	"Cycles." Stay with a technology for a finite number of periods, then continually switch to the frontier technology at each date for a finite number of periods, then stay with a technology for a finite number of

periods, etc.

Case IIA — Use inferior technology for $\tau \geq 0$ periods then "switch" to frontier technology at each subsequent date.

Case IV — Agent always "switches" to the frontier technology at each date.

Discussion of the Results.

Cases I and IV are straightforward. In case one, new vintage technologies do not represent an improvement that is big enough to offset the loss of informational human capital that switching to them entails, and therefore there is no switching, and no growth. In case IV, exactly the reverse is true.

Cases II and III, on the other hand, are a mix of cases I and IV in the above sense. At some levels of human capital, a switch pays, and for others it does not. Case II obtains when it is the human capital-rich agent that will switch, while in case III it is the human capital-poor agent that will do so. The latter is perhaps counter-intuitive in that human capital is usually thought of as being conducive to growth.

The most complicated dynamics arise in case IIB where cycles occur. We have illustrated a two-period cycle in figure 6, but we have not ruled out other more complicated ones. The most interesting, substantively, is probably case IIIA. Here, small differences in initial beliefs can have huge effects on long run outcomes, and perpetually widening inequality can arise among agents that are initially not too dissimilar.

V. Myopically Optimal Time Paths in the "Full-Menu" Model

In section III we remarked that there were three possible switching constraints one could impose. In the previous section we analyzed the first of those switching constraints - the NO JUMP model, where the agent may only choose between vintage between n and n+1. The NO-JUMP model is the most restrictive of the three switching constraints defined in section 3. We now study the "full-menu" model where at each date the agent may choose between a vintage n and vintages n+j for j=1,2,3,.... and where j is referred to as the jump size. All technologies are therefore always available, and the agent is aware of their existence, even though he may understand most of them extremely poorly in the sense that he may know virtually nothing about the relevant θ's. The full-menu model is the least restrictive of the switching constraints of section III.

Terminology and Assumptions. Recall the definition of h_2 in section III. Fix any date and refer to the status quo vintage at that date as vintage 0. If x is the posterior variance on the status quo vintage then $h_2(x)$ is the prior variance on vintage 1. The k-th iterate of h_2 from x, i.e., $h_2^k(x)$, gives the prior variance on vintage k.

Define G(x,k) to be the return from initial posterior variance x when a jump of size k is chosen. Note that G(x,0) is the return to choosing the status quo vintage and is equal to what was earlier referred to as $\ell(x)$. G(x,1) is the return to switching to the next vintage, with a jump of one; it is equal to what was earlier referred to as r(x). It should be easy to verify that

$$\begin{aligned} G(x,k) &= \gamma^k[1-\sigma_w^2-h_2^k(x)] \text{ and} \\ &= \gamma^k[1-\sigma_w^2-\alpha^k x-h_2^k(0)] \end{aligned}$$

where $h_2^k(0)$ is the k-th iterate of h_2 from 0 and where $h_2^0(x)=x$.

To make the problem interesting we shall impose the following assumptions:

A.1 $1-\sigma_w^2>0$.

A.2. $1-\sigma_w^2-\hat{x}<0$.

If A.1 is violated the return in each period, regardless of the jump size, is negative. The sum of expected discounted returns will then be negative regardless of the strategy chosen. A.1 rules this out. Condition A.2 states that as the agent upgrades his technology, his human capital depreciates fast enough to prevent him from attaining an arbitrarily large net output by choosing a large jump size. Indeed, suppose that A.2 is violated and in particular that $1-\sigma_w^2-\hat{x}>0$. Since $\hat{x}$ is the fixed point of the map h_2, from $\hat{x}$ iterates of h_2 remain at $\hat{x}$. The return to a jump of size k from initial posterior $\hat{x}$ will be $\gamma^k[1-\sigma_w^2-h_2^k(\hat{x})]=\gamma^k[1-\sigma_w^2-\hat{x}] \to \infty$ as $k\to\infty$. Hence when A.2 is violated it is possible to choose a jump size to obtain arbitrarily high return. Even the myopic problem is not well-defined that case. Condition A.2, together with A.1, ensure that the myopic problem is well-defined. (We will later on impose additional conditions to ensure that the infinite-horizon problem is also well defined.)

From assumption A.1, $G(0,0)\equiv 1-\sigma_w^2>0$ so there is at least one integer k such that $G(0,k)>0$. Now $h_2^k(0)$ converges monotonically from below to $\hat{x}$ as $k\to\infty$. Hence from A.2, there exists a unique $K<\infty$ such that $1-\sigma_w^2-h_2^k(0)>0$ for all $k\leq K$, and $1-\sigma_w^2-h_2^k(0)\leq 0$ for all $k>K$. In particular, $G(0,k)$ is positive for $k\leq K$ and is non-positive for $k>K$. Since $h_2^k(0)$ is monotone non-decreasing in k,

$$1-\sigma_w^2-h_2^K(0)>0\geq 1-\sigma_w^2-h_2^{K+1}(0)>\ldots>1-\sigma_w^2-h_2^{K+j}(0)>1-\sigma_w^2-h_2^{K+j+1}(0)>\ldots.$$

Since $\gamma>1$ this in turn implies that

$$0 \geq G(0,K+1) > \ldots > G(0,K+j) > G(0,K+j+1) > \ldots \tag{11}$$

For each $k=1,2,\ldots,$ $G(x, k-1)$ and $G(x, k)$ are linear in x and have slopes which are respectively $(\alpha\gamma)^{k-1}$ and $(\alpha\gamma)^k$. Hence so long as $\alpha\gamma\neq 1$, the two functions will have an intersection point (which may be negative). Define x_k^* to be the point of intersection of the two functions of x, G(x,k-1) and G(x,k); i.e.,

$$x_k^* \equiv \text{the unique } x \text{ such that } G(x,k\text{-}1)=G(x,k).$$

Notice that x_1^* is what in the NO JUMP model we referred to as x^*.

The Optimal Myopic Policy.

Recall that CASES I-IV were defined in section IV. The optimal policy functions for these cases in the NO JUMP model are as in fig. 4. We now provide, in the propositions below and in figures 7-9, the equivalent in the full-menu model. The reader may find it useful to compare fig. 4 with the corresponding figures 7-9. (As in section IV we ignore the knife-edge situations where either $\alpha\gamma=1$ or $r(0)=\ell(0)$.)

Proposition 5.1. (CASE I). Suppose that we are in CASE I. Then the payoff functions are as in fig. 7. In particular, $G(0,k)>G(0,k+1)$ for all k. Hence the optimal myopic action is to choose a jump size of zero (i.e., NO SWITCH) at each date from each initial posterior variance.

Proposition 5.2. (CASE II). Suppose that we are in CASE II. Then the payoff functions are as in fig. 8. In particular, there exists an integer $M \in \{1,\ldots,K\}$ such that

i. $G(0,0)<G(0,1)<\ldots<G(0,M)$; and

ii. $G(0,k)\leq G(0,M)$ for all $k>M$.

iii. The following jump sizes are optimal for the myopic problem from initial posterior variance $x\geq 0$:

- **a.** choose jump size 0 (i.e., action NO SWITCH) for x in $[x_1,\infty)$;
- **b.** jump size m if $x\in[x_{m+1}^*,x_m^*]$ for some $m=1,2,\ldots,M\text{-}1$; and
- **c.** jump size M if $x\in[0,x_M]$. So,

iv. The agent will never choose a jump size greater than M.

Proposition 5.3. (CASE III). Suppose that we are in CASE III. Then the payoff functions are as in fig. 9. In particular,

i. $G(0,k) > G(0,k+1)$ for all $k \geq 0$; and

ii. $\hat{x} < x_1^* < x_2^* < \ldots$, and $\lim_{k\to\infty} x_k^* = \infty$.

Hence,

iii. the optimal myopic action from any $x \geq 0$ is to choose a jump size $k=0,1,2,\ldots$, where k is any integer such that $x \in [x_k^*, x_{k+1}^*]$ (where $x_0^* \equiv 0$); and

iv. for any x and x′ with $x \geq x' \geq 0$, the optimal myopic jump size from initial posterior variance x is no less than the optimal myopic jump size from initial posterior variance x′.

Proposition 5.4. (CASE IV). Under assumption A.1 CASE IV can not occur.

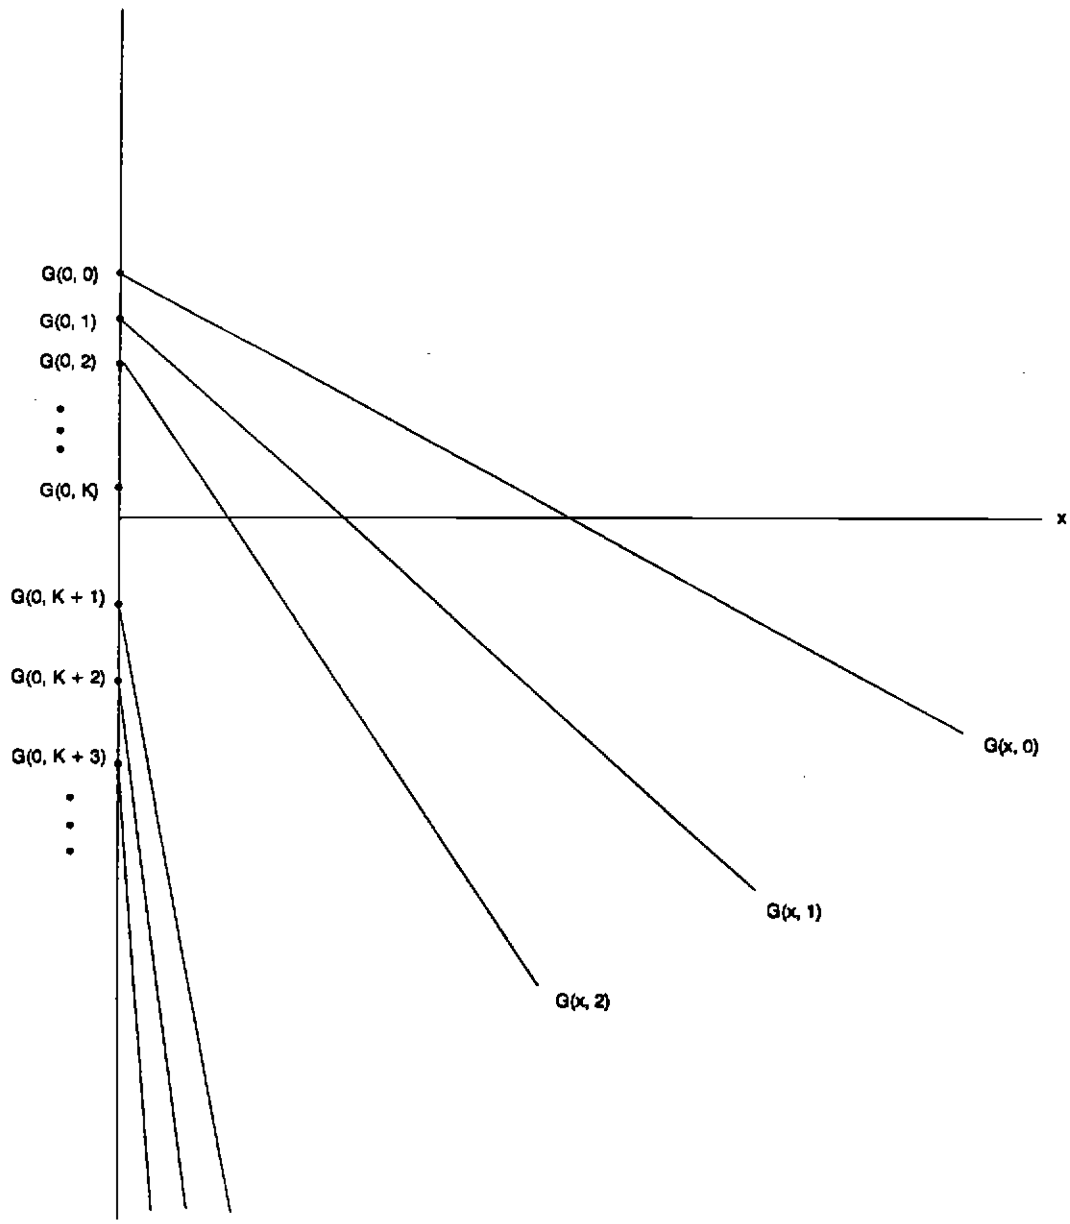

Figure 7: Full Menu Model in Case I.

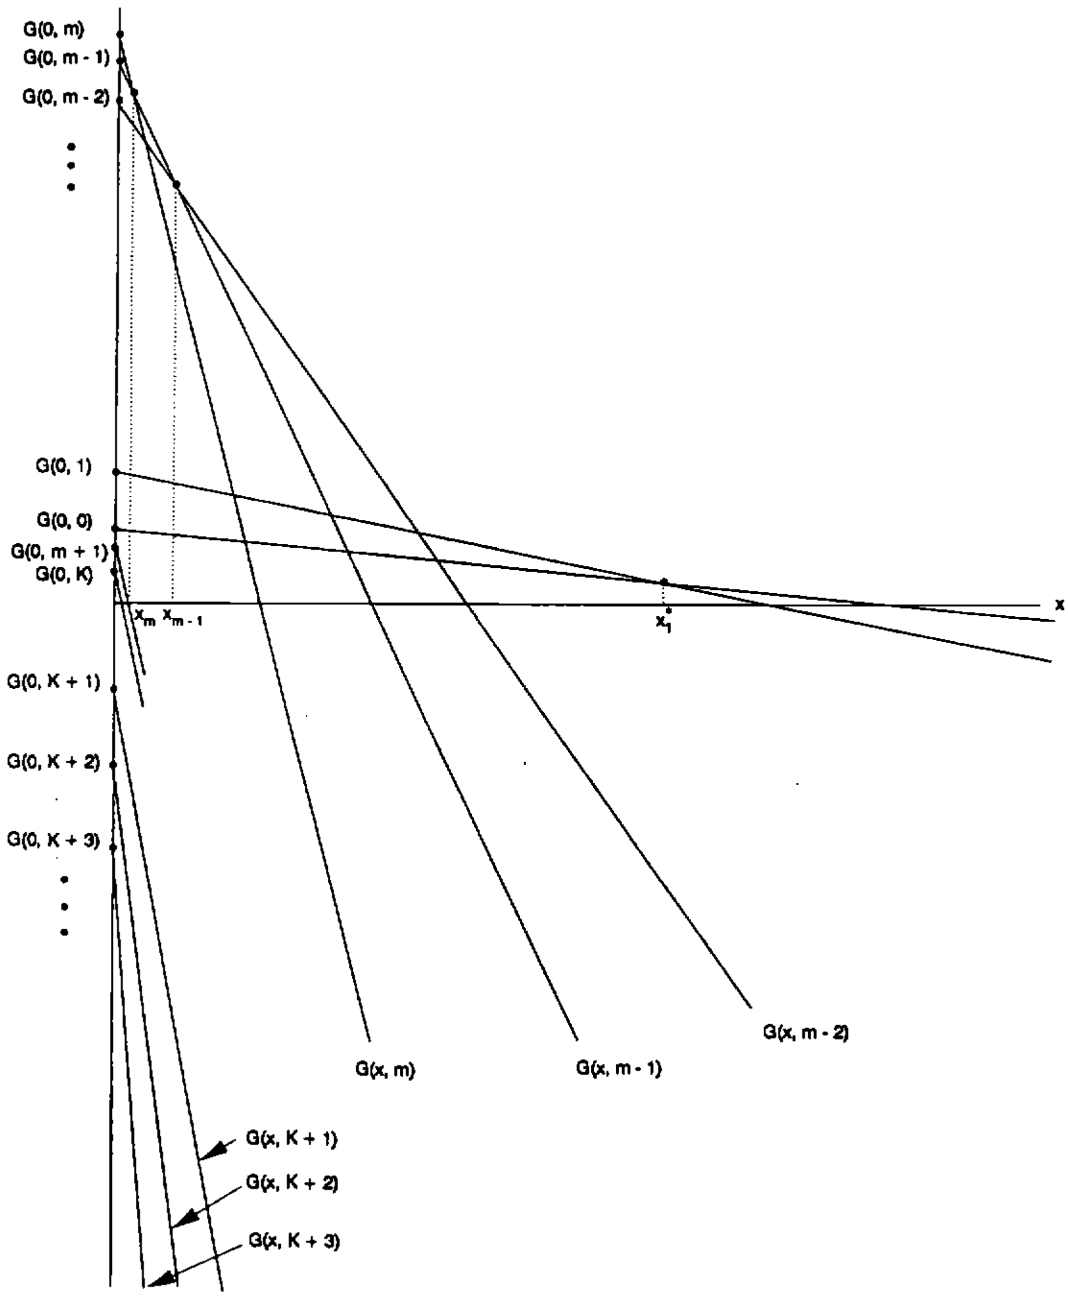

Figure 8: Full Menu Model in Case II.

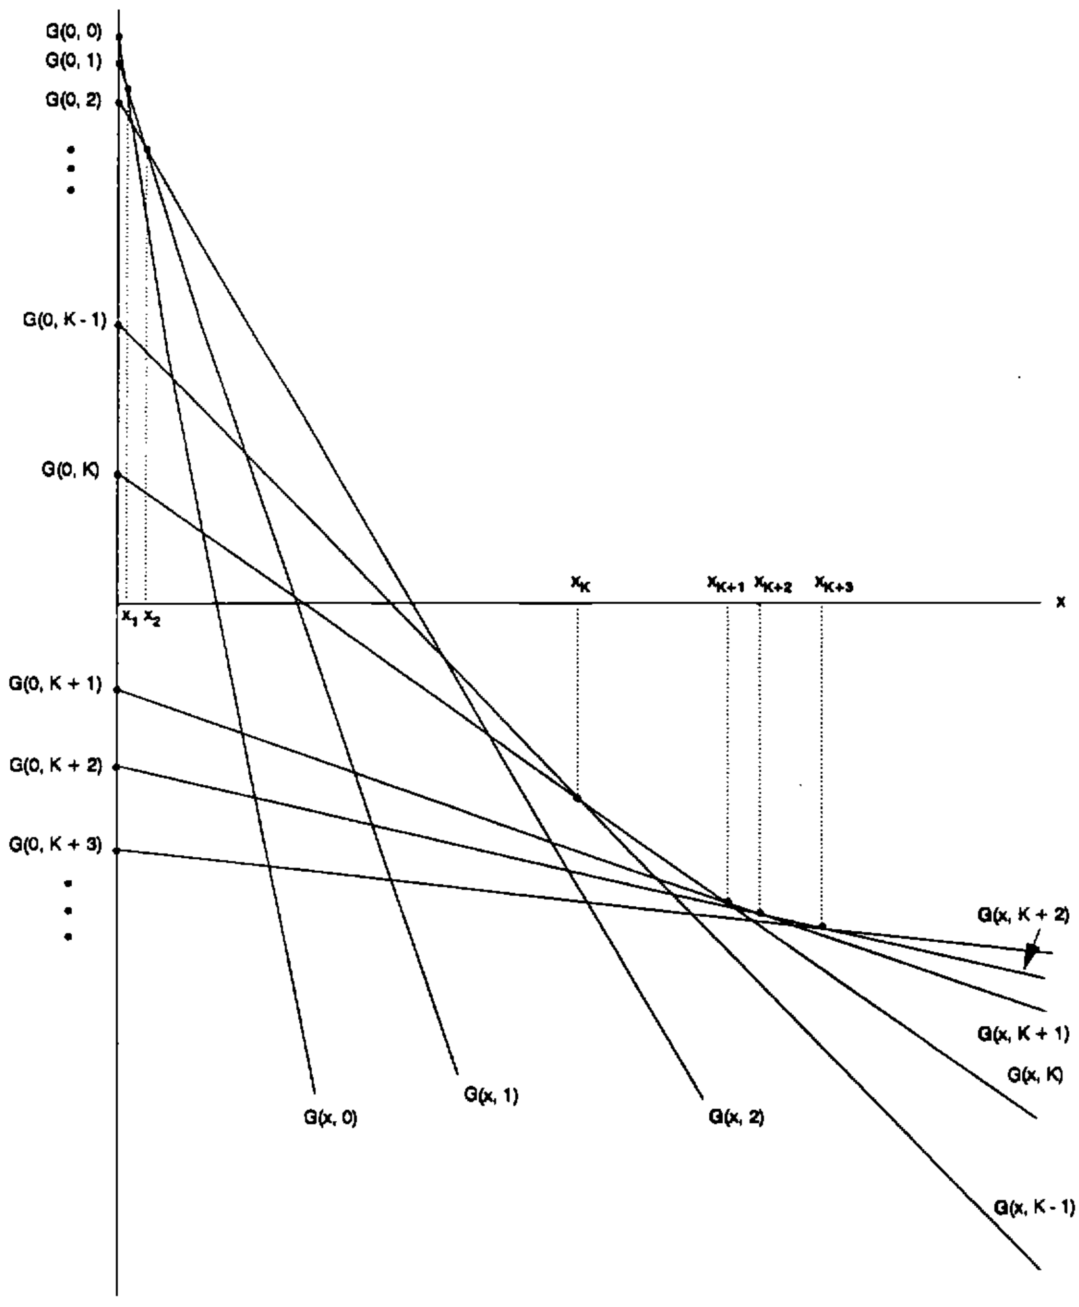

Figure 9: <u>Full Menu Model in Case III</u>.

Dynamics in the Myopic Full-Menu Model

Define for $k \geq 1$,

$$g_k(x) \equiv h_1(h_2^k(x)).$$

An agent with initial posterior variance x who chooses a jump size k will after using the new technology have a posterior variance of $g_k(x)$ over the new status quo technology. The map h_1 is concave while h_2^k is linear so $g_k(x)$ is concave in x. It is clearly increasing in x. Further, $g_k(0) > 0$. Further, since $h_1(x) \leq \sigma_w^2$, $g_k(x) \leq \sigma_w^2$ for all x. Hence $g_k(x)$ has a unique positive fixed point which we shall refer to as x_k^{**}. Note that what we referred to in section IV as x^{**} is what we are now referring to as x_1^{**}.

Lemma 5.6: **i.** $x_k^{**} < x_{k+1}^{**}$ for all $k \geq 1$; and **ii.** $\lim_{k\to\infty} x_k^{**} = h_1(\hat{x})$.

Corollary 5.7.: From any $x > h_1(\hat{x})$ the posterior variance of the status quo technology in the next period will be less than x.

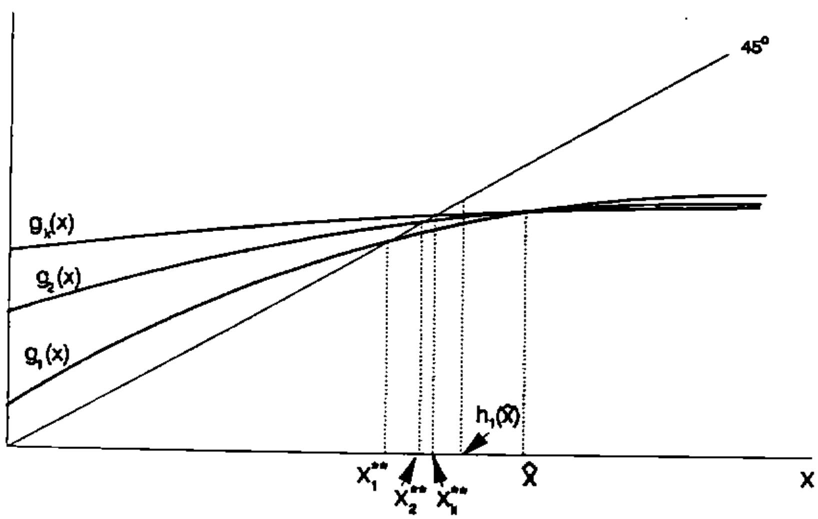

Figure 10: The g_k Functions.

We now describe the dynamics in the full-menu model using the classification of CASES that was used in the NO JUMP model. Our results will be very similar to the results in the analogous cases. We state our results in a collection of propositions. The proofs should be obvious from the figure (fig. 7-9) of the optimal policy corresponding to the case.

Proposition 5.8. (CASE I). Suppose that we are in CASE I. Then the optimal myopic action is to choose a jump size of zero (i.e., NO SWITCH) at each date from each initial posterior variance.

Proof: Follows immediately from fig. 7. ∎

CASE IIA': $x_1^{**} < x_1^{*}$ (**"eventually always choose positive jump size"**)

Recall that what we referred to in the NO JUMP MODEL as x^* (resp. x^{**}) is what we now refer to as x_1^* (resp. x_1^{**}). In the NO JUMP model we referred to CASE IIA as the case where $x_1^{**} < x_1^*$. In the NO JUMP model we showed that in this case eventually the agent will always choose to SWITCH. We now impose a condition slightly stronger than this. In particular, let M be as in Proposition 5.2 (or fig. 8); we assume that $x_M^{**} < x_1^*$. If $M=1$ then this condition is the same as what was previously called CASE IIA. It is a stronger condition when $M>1$. We shall call this stronger condition CASE IIA' to distinguish it from the previous condition in section IV. Under this condition we shall show that the agent will (for all but perhaps finitely many periods) choose in each period a jump size greater than or equal to one. Hence we obtain a result in the same spirit as the analogous case in the NO JUMP model. However, whereas in the NO JUMP MODEL eventually the jump size is exactly one, here eventually the jump size is at each date greater than or equal to one and may vary over time. Formally we have the following:

Proposition 5.9. (**CASE IIA'**): Suppose that we are in CASE IIA' (i.e., $x_M^{**} < x_1^*$). Then for all but possibly finitely many periods the agent will choose at each date t a jump size which is greater than or equal to one (and may depend on the date). In particular for all but perhaps finitely many periods, the action NO SWITCH will <u>not</u> be chosen.

CASE IIB: $x_1^* < x_1^{**}$ ("Cycles")

What was in the NO JUMP model referred to as CASE IIB is the situation where $x_1^* < x_1^{**}$. In the NO JUMP model the dynamics were characterized by "cycles." We now show that under the same condition we obtain cycles in the Full menu model as well. The only difference is that whereas in the NO JUMP model the cycles are between NO SWITCH and jump size 1, the cycles in this case will be between NO SWITCH and jumps of sizes $j \geq 1$.

Proposition 5.10. (Case IIB): Suppose that we are in CASE IIB ($x_1^* < x_1^{**}$). Then the dynamics are characterized by cycles: In particular, the action NO SWITCH is chosen for some time, then the agent decides to choose positive jumps for a while (with possibly time-varying jump sizes greater than or equal to one) then NO SWITCH is again chosen for a while then positive jumps occurs for a while, etc.

CASE IIIA: $x_1^* < x_1^{**}$ (Beliefs affect long run). We now impose the condition used in defining case IIIA in the NO JUMP model: $x_1^* < x_1^{**}$. In the NO JUMP model we obtained the conclusion that beliefs matter catastrophically. In particular, in the NO JUMP model for initial posterior variance x in $[0, x_1^*]$ it is optimal to choose the action NO SWITCH at each date while for x in (x_1^*, ∞) then it is optimal to SWITCH (or jump size one) at each date. Under the same condition we shall show that in the FULL MENU model a very similar result is obtained: for initial posterior variance x in $[0, x_1^*]$ it is optimal to choose the action NO SWITCH at each date; and for x in (x_1^*, ∞) it is optimal to choose a jump size, j, greater than or equal to one at each date.

Proposition 5.11. (CASE IIIA): Suppose that we are in CASE IIIA. **(i)** If the initial posterior variance lies in $[0, x_1^*]$ the optimal action is NO SWITCH (or jump size of zero) at <u>each date.</u>

(ii) If the initial posterior variance lies in (x_1^*, ∞) the optimal action is a jump size of one or greater <u>at each date</u> (with the size of jump possibly dependent on the date).

CASE IIIB′: $h_1(\hat{x}) < x_1^*$ (Switch for τ periods then "stick" with a vintage).

Recall that CASE IIIB in the NO JUMP model was the condition that $x_1^{**} < x_1^*$. We now define a stronger condition and refer to this as CASE IIIB′: $h_1(\hat{x}) < x_1^*$. Recall from lemma 5.6 that x_k^{**} converges from below to $x_\infty \equiv h_1(\hat{x})$. Hence case IIIB′ is the condition that $\lim_{k\to\infty} x_k^{**} < x_1^*$ and is in general stronger than CASE IIIB, $x_1^{**} < x_1^*$. For the NO JUMP model in case IIIB we showed that for all but possibly finitely many periods the agent will choose the action NO SWITCH. The posterior variance therefore converges to zero. We shall show that in the full-menu model the same result is true under the stronger condition of CASE IIIB′.

Proposition 5.12. (CASE IIIB′): Suppose we are in CASE IIIB′ (i.e., $h_1(\hat{x}) < x_1^*$).

(i) From (x_1^*, ∞) the posterior variance of the agent will enter the set $[0, x_1^*]$ in finite time.

(ii) Once the posterior variance is in the set $[0, x_1^*]$ the optimal action is to choose jump size 0 (or action NO SWITCH) in each and every period.

VI. Dynamically Optimal Policies in the No-Jump Model.

We now treat the case where the future matters, and in particular the agent's discount factor is positive. We will begin with the No-Jump model. We show that optimal policies exist, and are stationary in two senses: First, they do not depend on calendar time, and second, whether or not the agent holds on to his current (status quo) technology or switches to the next vintage does not depend on the numerical value of the vintage. In other words, whether or not the agent stays with vintage n or switches to n+1 depends only on his beliefs about vintage n, and is independent of the numerical value of n. We then characterize the dynamics under the positive discount factor assumption.

a. The Basic Structure and Notation.

At any date t the agent will have a history of previous vintages chosen and the observations associated with their use. A policy is a sequence $\pi = \{\pi_t\}_{t=1}^{\infty}$ where π_t maps the initial date 0 variance and the observed history at t into a date t decision SWITCH or NO SWITCH.

Fix a policy π. Let m(t) be the vintage chosen at t under π. At the start of date t+1 the agent chooses between the status quo vintage m(t) and the frontier vintage m(t) + 1. If the policy prescribes the action SWITCH at date t+1 then m(t+1) = m(t) + 1; and if the action is NO SWITCH then m(t+1) = m(t). The agent has a discount factor δ in (0,1). From initial posterior variance x, a policy π generates a sum of discounted expected payoffs given by

$$V_\pi(x) \;=\; \Sigma_{t=1}^{\infty}\delta^{t-1}\gamma^{m(t)}(1 - \mathrm{Var}_t\,\theta_{m(t)} - \sigma_w^2)\,, \tag{12}$$

where

$$\mathrm{Var}_t\,\theta_{m(t)} \;=\; E[(\theta_{m(t)} - E_t\theta_{m(t)})^2 \;\mid\; x,\,\pi]$$

is the variance of $\theta_{m(t)}$ under the agent's initial belief and policy π, which has expectations operator given by $E[\cdot \mid x,\pi]$ with initial posterior variance x over θ_1. We let Ψ denote the set of all policies. A policy π is said to be optimal if it attains supremum of $\mathrm{Sup}_{\pi' \in \Psi} V_{\pi'}(x)$ for all $x \geq 0$.

From initial posterior variance $x = x^{**}$, the strategy of switching in every period results in a sum of discounted net outputs equal to $r(x^{**})/[1-\gamma\delta]$ whenever $\delta < 1/\gamma$ and is infinite whenever $\delta \geq 1/\gamma$. To ensure that the sum of discounted payoffs in (12) is bounded, we impose the following

bound on the discount factor: $\delta < 1/\gamma$.

Implicit in the Bayes' rule mapping is the following bound on the variances over time. Fix any initial posterior variance over the status quo vintage in the first period. If the agent decides NO SWITCH in the first period then the posterior variance of the status quo vintage becomes $h_1(x)$. If the agent decides to SWITCH then the status quo vintage becomes vintage 2 and will have a posterior variance given by $h(x)$. More generally, the future posterior variances of the status quo vintage at that date will be given by iterates of the maps $h_1(\cdot)$ and $h(\cdot)$. The precise order and number of the iterates depends upon the decisions, SWITCH or NO SWITCH, chosen by the agent. Since $h_1(x) \leq x$ for all x, iterates of the map h_1 cause the posterior variance over the status quo vintage to fall. We defined earlier x^{**} to be the fixed point of the h mapping and indicated that iterates of this map converge monotonically to x^{**}. If x_1 is the date 1 posterior variance over the date 1 status quo technology, vintage 1, then it should be clear that any number of iterates of h and h_1 from x_1 will be bounded above by the maximum of x_1 and x^{**}. In particular, we obtain the following bound on the posterior variances over time:

$$\mathrm{Var}_t\, \theta_{m(t)} \leq \mathrm{Max}\, \{x_1, x^{**}\} \text{ for all } t.$$

Next we show that optimal policies exist by casting our model into a standard dynamic programming framework. The action space is {SWITCH, NO SWITCH}. Under the no recall assumption stated at the end of section III, the state variable is the pair (n, x) made up of the vintage, n, of the current status quo technology and the posterior variance x over the parameter θ_n of the status quo technology. The transition function from the date t state variable (n, x) space and the action "SWITCH" or "NO SWITCH" into the date $t+1$ state variable - either $(n+1, h(x))$ or $(n, h_1(x))$ - is continuous since h_1 and h are continuous. Hence standard dynamic programming arguments ensure the existence of an optimal policy under our boundedness assumptions.

When technology n is the status quo technology, the payoffs to any policy are proportional to γ^n. This means that the optimal policies will not depend on n, and that the value of the optimized objective will be proportional to γ^n. This is the content of the next lemma.

Let $V_n(x)$ be the value, under the optimal policy, of the sum of discounted payoffs when the status quo vintage is vintage n.

Proposition 6.1 (Stationarity)

(i) The optimal policy may be chosen to be independent of the numerical value of the status quo vintage and to depend only upon the variance of the status quo vintage.

(ii) $\qquad V_{n+1}(x) = \gamma V_n(x).$

Proof: Obvious.

Under this proposition, $V_n(x) = \gamma^{n-1}V_1(x)$ for all n. Let $V(x)$ denote $V_1(x)$, the value function in the "first" period, when the status quo technology is vintage one. Define

$$L(x) \equiv \ell(x) + \delta V(h_1(x)) \quad \text{and} \quad R(x) \equiv r(x) + \delta\gamma V(h(x)).$$

$L(x)$ (resp. $R(x)$) is the sum of discounted payoffs when the initial posterior variance over vintage 1 is x, when the action NO SWITCH (resp. SWITCH) is taken at date 1 and from date 2 onwards the optimal decision is chosen. The functional or Bellman equation is:

$$V(x) = \text{Max}\ \{L(x),\ R(x)\}.$$

We now have the following:

Proposition 6.2 (Convexity). $L(x)$, $R(x)$ and $V(x)$ are each downward-sloping, convex and continuous in x.

b. Dynamic Paths for Small x and for Large x.

In cases II and IV, for low values of x it is myopically optimal to switch, as shown in Figure 4. The proposition that one may have hoped to prove is "whenever it is myopically optimal to switch (i.e., whenever $r(x) \geq \ell(x)$) then it is also ***dynamically*** optimal to switch." The next proposition proves a somewhat weaker claim, but one that is in the same spirit.

Proposition 6.3: Fix any initial posterior variance x and suppose that $r(x) \geq \ell(x)$ and $x \leq \hat{x}$. Then from x the strategy of choosing the action NO SWITCH at ***each date*** is ***not*** optimal. In particular there exists an $x' \in (0,x)$ such that from x' the optimal action is SWITCH.

All that this proposition is ruling out is the possibility that the agent chooses action NO SWITCH at ***each date.*** It however allows for NO SWITCH at some date so long as the agent also chooses SWITCH at some other date. A stronger claim however is true when $x=0$:

Corollary 6.4. Suppose that $r(0) \geq \ell(0)$. (Note that this occurs in CASES II and IV.) Then the optimal action from $x=0$ is SWITCH.

Remark: It should be clear from the proof of Corollary 6.4 that the corollary can be made to hold even when $r(0) < \ell(0)$ so long as $r(0)$ is not "too much smaller than" $\ell(0)$. By continuity, the action SWITCH is taken in a neighborhood of $x = 0$.

The previous two results provide some insight into the dynamics for small values of x, the initial posterior. We now say something about the large x situation. The next proposition follows from a property illustrated in figure 1: h_1 and h are both bounded above by σ_w^2. This means that the informational value of current decisions is bounded. And since the current payoffs are unbounded as x gets large, this means that for large enough x, current considerations will dominate relative to future ones, and they alone will determine the optimal action. This is the content of the next proposition, although it should be taken with some restraint since for large x the value function is negative so perhaps production will not even take place!

Proposition 6.5: Fix a discount factor $\delta>0$. There exists an $\bar{x}>0$ sufficiently large such that for all $x>\bar{x}$ the optimal action for the infinite horizon $\delta>0$ problem is the same as that of the myopic problem.

c. Dynamic Paths for Large δ

When the discount factor gets large, the critical feature of the model becomes the value of $r(x^{**})$. Since x^{**} is the limit of x for an agent who switches technologies in every period, this means that if an agent switches at the maximal rate, this policy will sustain a positive net output only if $r(x^{**})$ is positive. In this case, a policy that maximizes the rate of growth will indeed be optimal for large enough discount factors, and this is the content of proposition 6.6.

But when $r(x^{**})$ is negative, switching at the maximal rate yields negative long run net output, and can not be optimal. To sustain positive long-run output, the agent must pause after switching, at least occasionally. Corollary 6.4 implies that it may not be optimal to pause forever, in which case there will be cycles (in the same sense as case IIB), and this is the content of proposition 6.8.

It is straightforward (but tedious) to show that

$$r(x^{**}) \begin{array}{c}\geq\\=\\\leq\end{array} 0 \qquad \text{as} \qquad \sigma_\epsilon^2 \begin{array}{c}\leq\\=\\\geq\end{array} 1 - (1+\alpha)\sigma_w^2 + \alpha\sigma_w^2.$$

In particular, $r(x^{**})$ can be either positive or negative depending upon the values of α, σ_ϵ^2 and σ_w^2.

Proposition 6.6: Assume that $r(x^{**}) > 0$. Fix any $\bar{\sigma}^2 < \infty$. Then there exists a $\bar{\delta}$ in $(0, 1/\gamma)$ such that for all $\delta \in (\bar{\delta}, 1/\gamma)$, the optimal action from any initial posterior variance x in $[0, \bar{\sigma}^2]$ is SWITCH.

The restriction that $r(x^{**}) \geq 0$ can in fact coexist with all six cases: I, IIA, IIB, IIIA, IIIB, and IV. That is, the set of parameters for which $r(x^{**}) \geq 0$ has a non-empty intersection with each of the six subsets of the parameter space. Appendix 2 reports six combinations of parameters, one for each of the six cases, each of which satisfies the restriction $r(x^{**}) \geq 0$. This means that when $r(x^{**})$ is positive, optimal policies behave very differently when δ is large compared to how they look when δ is small or zero. When δ is small, continual switching is optimal only when we are in case IV, whereas when δ is large, switching is optimal for all parameters satisfying $r(x^{**}) \geq 0$.

One of the differences between the myopic model and the dynamically optimal (i.e. positive discount factor) model is the following. In the myopic model we have:

Proposition 6.7: For the <u>myopic</u> model the agent's expected output is an increasing function of time.

The proof of this proposition is the following: Suppose that at date t the agent chooses a vintage n and obtains an expected output level of Y_t. At date t+1 if the same vintage n is chosen, since the agent will have better information on it, the output level from the use of vintage n at date t+1, Y'_{t+1} say, will be higher than Y_t. For the myopic model, at date t+1 the agent will only choose a vintage different from vintage n if that vintage yields a higher expected output. This means that the date t+1 output level can not be less than Y'_{t+1}. Combining these arguments shows that $Y_{t+1} \geq Y_t$. In summary, for the myopic model an agent either sticks to the same vintage (which because of declining variances results in higher expected output) or switches to another vintage if that vintage results in a yet higher expected output. In either case expected outputs must rise over time.

This is not the case when the discount factor is positive. In particular, in the positive discount factor model an agent may sacrifice current output in order to attain experience with a better vintage. This is an important difference between the positive discount factor and zero discount factor models. We now illustrate this possibility with the aid of Proposition 6.6.

Suppose that $r(x^{**})>0$. Choose the discount factor sufficiently large so that Proposition 6.6 holds from all posterior variances $x \leq x^{**}$. In particular for such a discount factor it will be optimal for the agent to SWITCH at each and every date from any beginning of period posterior variance $x \leq x^{**}$. Consider an agent who begins with initial posterior variance on the status quo vintage, vintage 1, equal to zero. Under the optimal policy the dates one and two expected output levels will be

$$y_1 = \gamma[1-\sigma_w^2-\sigma_\epsilon^2] \quad \text{and} \quad y_2 = \gamma^2[1-\sigma_w^2-\alpha h_1(\sigma_\epsilon^2)-\sigma_\epsilon^2].$$

Define

$$\bar{\gamma} \equiv (1-\sigma_w^2-\sigma_\epsilon^2)/(1-\sigma_w^2-\alpha h_1(\sigma_\epsilon^2)-\sigma_\epsilon^2).$$

It is easy to verify that for $\gamma<\bar{\gamma}$, $y_1 > y_2$. In particular as long as the parameter γ is not too large, for this agent the optimal strategy will initially result in a <u>reduction</u> in output over time!

Of course these declines in output can not persist over time. Indeed, let Y_t denote the date t expected output and let x_t denote the date t posterior variance of vintage chosen at date

t. Then the difference in expected output levels is given by

$$[Y_{t+1} - Y_t]/\gamma^{t-1} = \gamma r(x_{t+1}) - r(x_t) \rightarrow (\gamma-1)r(x^{**}) > 0 \text{ as } t \rightarrow \infty.$$

Hence in the limit the expected output, Y_t, increases at the rate γ.

We provide a result which shows what may happen when $r(x^{**}) < 0$.

Proposition 6.8: Suppose that $r(0) > \ell(0)$ and $r(x^{**}) < 0$. Then there exists a $\bar{\delta}$ in $(0, 1/\gamma)$ such that for all $\delta \in (\bar{\delta}, 1/\gamma)$, from any initial posterior variance, the process of actions of agents will be characterized by cycles: The agent will choose NO SWITCH for a while, then SWITCH for a while, then NO SWITCH for a while, etc.

Proposition 6.8 assumes the restriction $r(0) > \ell(0)$, which rules out cases I and III. The additional restriction $r(x^{**}) < 0$ also rules out case IV. This last assertion is proved in Appendix 3. After an extensive computer search, we suspect that these restrictions also rule out case IIA, although we have not proved this claim. We do know, however, that the restrictions hold on a non-empty subset of parameters satisfying case IIB. A point in this subset is the vector $(\sigma_w^2 = 0.5, \sigma_\epsilon^2 = 0.28, \alpha = 0.9, \gamma = 2.3)$. Therefore, as far as we can tell, Proposition 6.8 applies only to a subset of case IIA, and here the outcome is "cycling" for large δ just as it was for small δ.

d. Dynamic Paths for Small δ

If the discount factor is strictly positive but small, one would expect that the dynamically optimal policies and time-paths will resemble the myopically optimal ones. This section shows that this is indeed so.

To avoid uninteresting details, in the proposition below we suppose that the parameter values $(\alpha, \gamma, \sigma_w^2, \sigma_\epsilon^2)$ are in the *generic* set

$$G \equiv \{(\alpha,\gamma,\sigma_w^2,\sigma_\epsilon^2): \ \gamma\alpha \neq 1 \text{ and } [1 - \sigma_w^2] \neq \gamma[1 - \sigma_\epsilon^2 - \sigma_w^2]\}.$$

Let $a_\delta^*(x)$ denote the optimal current period action, either SWITCH or NO SWITCH, for the infinite horizon problem with discount factor $\delta > 0$ when the initial posterior variance over the current status quo technology is x. Let $a_0^*(x)$ denote the optimal current period action for the myopic problem from the initial posterior variance x.

Proposition 6.9. (Small δ).

(i) Suppose that we are in CASE I or CASE IV. Then there exists a $\bar{\delta} > 0$ such that for all δ in $[0, \bar{\delta})$ and for all $x \geq 0$, $a_\delta^*(x) = a_0^*(x)$; i.e., the optimal action from beginning of period posterior variance x with discount factor δ is the same as the optimal current period action in the myopic problem with beginning of period posterior variance x.

(ii) Suppose that we are in either CASE II or CASE III. Then for all $\xi > 0$ there exists a $\bar{\delta} > 0$ such that for all δ in $[0, \bar{\delta})$ and for all $x \geq 0$ outside of the neighborhood $(x^* - \xi, x^* + \xi)$ of x^*, $a_\delta^*(x) = a_0^*(x)$.

The proposition above merely states the optimal action in the <u>current</u> period for the small discount factor problem is the same as that of the myopic or zero discount factor problem. It does not say what happens over time. The proposition is true only for a set of initial posterior variances in a subset, $[0, x^* - \xi] \cup [x^* + \xi, \infty)$, of the real line. If the dynamics of the problem never cause the posterior variance of any vintage at any date to enter the set $(x^* - \xi, x^* + \xi)$ then we may conclude that the entire dynamic process for the small discount factor problem is the same as that of the myopic problem. This is the case in the situations listed below. Hence in these situations the myopic model and the small discount factor problem have identical actions *at each and every date*. We omit the proofs as these may be easily verified from observing the figures associated with each of the cases mentioned.

Corollary 6.10 (Dynamics with small δ). Let x denote the initial date one posterior variance over the initial technology, vintage 1. Fix any $\xi > 0$. Then there exists a $\bar{\delta} > 0$ such that for all δ in $[0, \bar{\delta})$ the infinite horizon problem with discount factor δ and the myopic problem result in the same optimal decisions *at each date* in each of the following situations:

(I) In CASE I or IV for all $x \geq 0$;

(IIA) In CASE IIA for all $x \leq x^* - \xi$. However, even when $x > x^* - \xi$, the myopic problem and the small discount factor can differ in optimal actions for only *finitely* many periods. The *long run* behavior of the two models is the same (and in particular in all periods except for possibly finitely many initial periods, the agent will SWITCH in each period).

(IIIA) In Case IIIA for all x in $[0, x^* - \xi] \cup [x^* + \xi, \infty)$ (where we suppose ξ may be chosen sufficiently small so that $x^{**} > x^* + \xi$).

(IIIB) In CASE IIIB for all $x \leq x^* - \xi$. (However, even when $x > x^* - \xi$, the myopic problem and the small discount factor can differ in optimal actions for only *finitely* many periods. The long run behavior of the two models are the same. In particular in all periods except for possibly finitely many initial periods, the agent will choose NO SWITCH in each period).

The only situation where the corollary above does not apply[4] is in CASE IIB. Note that

---

[4] The corollary does not deal with the non-generic case, which may be handled as follows: Suppose that the parameter values are outside of the set G. Then either $\gamma\alpha = 1$ or $[1 - \sigma_w^2] = \gamma[1 - \sigma_\epsilon^2 - \sigma_w^2]$. We consider the three possibilities: (i) Suppose first that $\gamma\alpha = 1$ ***and*** $[1 - \sigma_w^2] = \gamma[1 - \sigma_\epsilon^2 - \sigma_w^2]$. Then $\ell(x) = r(x)$ for all $x \geq 0$, and so both actions SWITCH and NO SWITCH are optimal for the myopic problem. The conclusions of the corollary therefore hold trivially for all $x \geq 0$. (ii) Next suppose that $\gamma\alpha = 1$ and $[1 - \sigma_w^2] \neq \gamma[1 - \sigma_\epsilon^2 - \sigma_w^2]$. It is straightforward to check that the proof of Part (i) of Proposition 6.9 only used the requirement that $[1 - \sigma_w^2] \neq \gamma[1 - \sigma_\epsilon^2 - \sigma_w^2]$. So the corollary above holds in this situation. (iii) Finally, suppose that $\gamma\alpha \neq 1$ and $[1 - \sigma_w^2] = \gamma[1 - \sigma_\epsilon^2 - \sigma_w^2]$. The intersection of $r(x)$ and $\ell(x)$ is now $x^* = 0$. We therefore necessarily have $x^* \leq x^{**}$. Hence we are essentially in CASE IIB (if $\gamma\alpha > 1$) or CASE IIIA (if $\alpha\gamma < 1$), but with $x^* = 0$. Part (ii) of the Proposition then applies but for x in $(x^* + \xi, \infty)$. Also, and the analogous parts of the

the problem is that in the neighborhood of radius ξ around the intersection x^* of $\ell(x)$ and $r(x)$, the optimal current period actions of the myopic and the infinite horizon problem may differ. In case IIB it is in principal possible that from some given initial posterior variance $x > 0$, the subsequent posterior variances will visit that neighborhood infinitely often and hence in principle it is possible that the myopic and infinite horizon optimal actions may differ at infinitely many dates.

Remark. The one real difference we have noticed between the myopic case and the small $\delta > 0$ case is the following: Suppose we are in CASE III and suppose that $x^* < \hat{x}$. In the myopic model, from initial posterior variance equal to x^* it is optimal to choose the action NO SWITCH at each and every date. However, from Proposition 6.3 we know that so long as the discount factor is positive, however small that might be, it is ***not*** optimal to choose the action NO SWITCH at each and every date.

---

corollary then hold.

VII. Positive Discount Factor in Full-Menu Model

Boundedness of the Value Function:

Recall that from the definition of K in (11), $1-\sigma_w^2-h_2^{K+1}(0) \leq 0$. Define $\Psi: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ by

$$\Psi(x) \equiv \gamma^{K+1}(1-\sigma_w^2)+\delta\gamma^{K+1}x, \text{ and}$$

$$x_\Psi \equiv \gamma^{K+1}(1-\sigma_w^2)/[1-\delta\gamma^{K+1}].$$

We impose the following two assumptions:

A.3. $\delta\gamma^{K+1}<1$; and

A.4. $1-\sigma_w^2-h_2^{K+1}(0)+\delta\gamma^{K+1}x_\Psi<0$.

It should be clear that A.3 will hold when δ is small. Also since $1-\sigma_w^2-h_2^{K+1}(0)\leq 0$, A.4 will also hold for sufficiently small δ. Under A.3 the function $\Psi(x)$ is a linear function with slope less than one and with unique fixed point x_Ψ.

Define $V^T(x)$ to be the value function of the T-horizon problem when the current posterior variance is x and the status quo technology is vintage 0. Note that if instead the status quo vintage is some vintage k then the T-horizon value function becomes $\gamma^k V^T(x)$.

Lemma 7.1 (Bound on Value Function). $V^T(x) \leq x_\Psi$ for all T and for all x.

Since the value function for the infinite horizon model is equal to the limit as $T\rightarrow\infty$ of the T-horizon value functions, the above result implies the boundedness of the former.

Appendix 1. The Proofs.

Proof of Lemma 3.1.: We begin by proving a stronger version of (iii):

Claim 3.1.1. For all $n=1,2,\ldots,$ and for all $x\leq\hat{x}$, $h_1(h_2^n(x))<h_2^n(h_1(x))$.

Proof of Claim 3.1.1. Fix any $n=1,2,\ldots$, and define $\phi_{F,n}(x) \equiv h_1(h_2^n(x))$ and $\phi_{S,n}(x) \equiv h_2^n(h_1(x))$. It should be clear that $\phi_{F,n}(0)=h_1(h_2^n(0)) < h_2^n(0) = \phi_{S,n}(0)$. Also, taking derivatives we conclude that $\phi_{F,n}'(x)=\alpha^{n-1}h_1'(h_2^n(x))$ and $\phi_{S,n}'(x)=\alpha^{n-1}h_1'(x)$. For $x \leq \hat{x}$, $h_2^n(x) \geq x$. Since h_1 is strictly concave this implies that for $0 \leq x \leq \hat{x}$, $\phi_F'(x) \leq \phi_S'(x)$, with strict inequality for $x < \hat{x}$. This proves that for $x \leq \hat{x}$, $\phi_{F,n}(x) < \phi_{S,n}(x)$. ▮

Proof of Lemma 3.1 (cont'd): We shall prove this lemma in reverse order.

(iii) Part (iii) is the same as claim 3.1.1 above with $n=1$.

(ii) From part (iii) we know that condition LPSY holds for all $x \leq \hat{x}$. So it remains only to show that the condition holds at any $x > \hat{x}$. So fix such an x. Then by definition of $\hat{x}$, $x > h_2(x)$ so

$$h_1(x) > h_1(h_2(x)) \equiv \phi_F(x). \tag{13}$$

Since $h_1(x) \leq \sigma_w^2$, when $\hat{x} \geq \sigma_w^2$ we conclude that $h_1(x) \leq \hat{x}$; i.e., $h_1(x)$ lies to the "left" of the fixed point $\hat{x}$ of the h_2 map. Hence $h_1(x) \leq h_2(h_1(x)) \equiv \phi_S(x)$. (13) therefore implies that condition LPSY holds at x.

(i) It is easy to check that (with obvious abuse of notation) $\phi_S(\infty)=h_2(h_1(\infty))=h_2(\sigma_w^2)= \alpha\sigma_w^2+\sigma_\epsilon^2$ and $\phi_F(\infty)=h_1(h_2(\infty))=\sigma_w^2$. Hence under the hypothesis of this part of the lemma, $\alpha\sigma_w^2+\sigma_\epsilon^2<\sigma_w^2$ so $\phi_S(\infty)<\phi_F(\infty)$. Hence $x_{LPSY}<\infty$. Since $\hat{x}>0$ we conclude from part (iii) that $x_{LPSY}>0$. Hence $0<x_{LPSY}<\infty$.

Next, fix any $x \geq x_{LPSY}$. Now, $\phi_F'(x)=\alpha h_1'(h_2(x))$ and $\phi_S'(x)=\alpha h_1'(x)$. Part (iii) of this lemma implies that for $x > x_{LPSY}$, $x > \hat{x}$. This in turn implies that $x > h_2(x)$. Since h_1 is concave this implies that $\phi_F'(x) > \phi_S'(x)$ at each $x > x_{LPSY}$. Since $\phi_F(x_{LPSY})=\phi_S(x_{LPSY})$, we conclude that $\phi_F(x) > \phi_S(x)$ for all $x > x_{LPSY}$. ▮

Proof of Proposition 3.2. We proceed via two lemmas:

Lemma 3.2.1. Fix any $x < x_{LPSY}$. Then for each $n=1,2,\ldots$, $h_1(h_2^n(x)) < h_2^n(h_1(x))$.

Proof: From claim 3.1.1 in the proof of Lemma 3.1 we know that this lemma is true for $x \leq \hat{x}$. So fix any x in the interval $(\hat{x}, x_{LPSY})$. Since iterates of the h_2 map are decreasing for any $x > \hat{x}$, if $x < x_{LPSY}$ then so too is $h_2^n(x) < x_{LPSY}$. In particular condition LPSY holds at $h_2^n(x)$ so $\phi_F(h_2^n(x)) < \phi_S(h_2^n(x))$ or

$$h_1(h_2(h_2^n(x)) < h_2(h_1(h_2^n(x)).$$

We proceed by induction. The lemma is trivially true for $n=1$ since this is implied by the definition of x_{LPSY}. Suppose that the conclusion of this lemma holds for some n. In particular suppose that for some n, $h_1(h_2^n(x)) < h_2^n(h_1(x))$. Applying the h_2 operator to both sides of this inequality implies that

$$h_2(h_1(h_2^n(x))) < h_2(h_2^n(h_1(x))).$$

The previous two inequalities imply that $h_1(h_2^{n+1}(x)) < h_2^{n+1}(h_1(x))$, so the lemma holds for $n+1$. Hence by induction the lemma is true for all n. ∎

Lemma 3.2.2. Fix any two agents A and B. Suppose that the date 1 variance on vintage 1 obeys condition LPSY. Fix any date T and any vintage $\bar{n}$. Suppose that at the beginning of date T neither agent has sampled any technology of vintage $n \geq \bar{n}$. Let $\mathrm{Var}_A\, \theta_{\bar{n}}$ and $\mathrm{Var}_B\, \theta_{\bar{n}}$ denote the beginning of date T posterior variances over vintage $\bar{n}$ of the two agents A and B respectively. Suppose that $\mathrm{Var}_A\, \theta_{\bar{n}} \leq \mathrm{Var}_B\, \theta_{\bar{n}}$. Fix any two vintages n' and n" with $\bar{n} \leq n' \leq n''$. Suppose that at date T agent A samples vintage n" while agent B samples the "lower" vintage n'. Fix any $N \geq n''$. Let $\mathrm{Var}_A\, [\theta_N \mid \langle n'' \rangle]$ and $\mathrm{Var}_B\, [\theta_N \mid \langle n' \rangle]$ denote the posterior variances of agents A and B respectively over vintage N. Then

$$\mathrm{Var}_A\, [\theta_N \mid \langle n'' \rangle] \leq \mathrm{Var}_B\, [\theta_N \mid \langle n' \rangle],$$

with strict inequality if either $\mathrm{Var}_A\, \theta_{\bar{n}} < \mathrm{Var}_B\, \theta_{\bar{n}}$ or $n' < n''$.

Proof of Lemma 3.2.2.: It is easy to verify that

$\mathrm{Var}_A\ [\theta_n'' \mid <n''>] = h_1(h_2^{n''-\bar{n}}(\mathrm{Var}_A\ \theta_{\bar{n}})) = h_1(h_2^{n''-n'}(m_A))$ where $m_A = h_2^{n'-\bar{n}}(\mathrm{Var}_A\ \theta_{\bar{n}})$; and

$\mathrm{Var}_B\ [\theta_n'' \mid <n'>] = h_2^{n''-n'}(h_1(h_2^{n'-\bar{n}}(\mathrm{Var}_B\ \theta_{\bar{n}}))) = h_2^{n''-n'}(h_1(m_B))$ where $m_B = h_2^{n'-\bar{n}}(\mathrm{Var}_B\ \theta_{\bar{n}})$.

Now, $\mathrm{Var}_A\ \theta_{\bar{n}} \leq \mathrm{Var}_B\ \theta_{\bar{n}}$ implies that $m_A \leq m_B$. Hence we may apply lemma 3.2.1 to conclude the proof of this lemma. (The application of lemma 3.2.1 requires condition LPSY to hold for the date T variances. Note that since $[0,\hat{x}] \subseteq [0,x_{LPSY}]$, the latter interval is a stable set in the sense that if the date 1 variances of vintage 1 lie in this set, so too will all subsequent variances of all subsequent technologies. Since this lemma assumes date 1 variances obey condition LPSY, the same will be true of the date T variances.) ∎

Proof of Proposition 3.2 (cont'd): We shall prove this by induction on the date. The case for $T=1$ is handled by the Lemma 3.2.2. So assume that this proposition is true for some T for any individuals F and S satisfying the hypotheses of this proposition. Since the posterior variance at date $T+1$ of any agent is independent of the order in which the vintages were sampled, we may suppose that agents sample vintages in a non-decreasing order so that the vintage chosen at any date t is no less than that chosen at date t-1. Let $\bar{n}$ denote the vintage chosen at date T by agent F. Let n'' (resp. n') be the vintage chosen by agent F (resp. S) at date $T+1$. From the induction hypothesis, the date T variance over vintage $\bar{n}$ of agent F, $\mathrm{Var}_{F,T}\ \theta_{\bar{n}}$ will be less than or equal to that of agent S, $\mathrm{Var}_{S,T}\ \theta_{\bar{n}}$. Suppose first that $n' \geq \bar{n}$. Then from lemma 3.2.2 above, $\mathrm{Var}_{F,T+1}\ \theta_{n''} \leq \mathrm{Var}_{S,T+1}\ \theta_{n''}$, which is the induction hypothesis for $T+1$. Suppose instead that $n' < \bar{n}$. Consider another agent, who we shall call agent $\hat{S}$, who behaves like agent S in periods 1 through T, but at date $T+1$ chooses vintage $\bar{n}$. From the argument we just made above we know that for agent $\hat{S}$,

$$\mathrm{Var}_{F,T+1}\ \theta_{n''} \leq \mathrm{Var}_{\hat{S},T+1}\ \theta_{n''}.$$

However from Lemma 3.2.2 we may conclude that

$$\text{Var}_{\hat{S},T+1}\,\theta_{n^*} < \text{Var}_{S,T+1}\,\theta_{n^*}.$$

Hence $\text{Var}_{F,T+1}\,\theta_{n^*} < \text{Var}_{S,T+1}\,\theta_{n^*}$, which again is the induction hypothesis for $T+1$. Hence by induction the proposition is true for all T. (It should be clear how we obtain the strict inequality asserted in this proposition from the strict inequalities obtained in Lemma 3.2.2.) ∎

Proofs of Propositions 5.1-5.4: We begin with the following lemma:

Lemma 5.5. Assume $\alpha\gamma \neq 1$.

i. Then for each $k=1,2,\ldots$, $x_k^*=h_2^{-k+1}(x_1^*)$. In particular, x_k is the (k-1)-th iterate of the inverse of the function h_2 from x_1^*.

ii. If $x_k \leq 0$ for some k, then $x_{k'} < 0$ for all $k' > k$.

Proof: (i) By definition x_1^* is the unique point of intersection of the functions G(x,0) and G(x,1). Fix any integer k. From the definition of x_k, $G(x_k,k-1)=G(x_k,k)$. Hence from the definition of the G(x,k) functions, $[1-\sigma_w^2-h_2^{k-1}(x_k^*)]=\gamma[1-\sigma_w^2-h_2(h_2^k(x_k^*))]$. In particular, at $x=h_2^{k-1}(x_k^*)$ we have $G(x,0)=G(x,1)$. Hence $h_2^{k-1}(x_k^*)=x_1^*$, from which part (i) of this lemma follows.

(ii) From (i) we know that the x_k^*'s are iterates of the map h_2^{-1}. The slope of the latter is $1/\alpha$. When $\alpha<1$, $1/\alpha>1$ so iterates from any $x_k^* \leq 0$ converge monotonically to $-\infty$. If $\alpha=1$ the h_2^{-1} mapping has slope of one but is everywhere below the 45 degree line. Again, iterates from any $x_k^* \leq 0$ converge monotonically to $-\infty$. Finally, if $\alpha>1$, the fixed point of h_2^{-1} is the point $\sigma_\epsilon^2/(1-\alpha)$ which is negative. From $x_k^* \leq 0$, iterates of h_2^{-1} converge monotonically to this point. In either case, part (ii) of the lemma follows. ∎

Proof of Proposition 5.1. (CASE I). In case I $\alpha\gamma>1$. Also $G(0,0)=\ell(0) > r(0)=G(0,1)$. So $x_1^*<0$. From lemma 5.5(ii) this implies that $x_k^*<0$ for all k. If for some k $G(0,k) \leq G(0,k+1)$ then, since the absolute value of the slope of $G(x,k+1)$ exceeds that of

$G(x,k)$ when $\alpha\gamma > 1$, their point of intersection, x_{k+1}^*, will be positive. This is a contradiction. Hence $G(0,k) > G(0,k+1)$ for all k from which the proposition follows. ∎

Proof of Proposition 5.2. **(CASE II)**. In case II $\alpha\gamma > 1$ and $G(0,0) = \ell(0) < r(0) = G(0,1)$. Hence $x_1^* > 0$. Define M to be the largest integer m such that $G(0,k\text{-}1) < G(0,k)$ for $k = 1,2,\ldots,m$. From the definition of M, $G(0,M) \geq G(0,M+1)$. Since the absolute values of the slopes of the $G(x,k)$ functions are increasing in k when $\alpha\gamma > 1$, this implies that $x_{M+1}^* \leq 0$. From lemma 5.5(ii) this in turn implies that $x_k^* \leq 0$ for all $k > M$. Suppose, per absurdem, that for some $k > M$, $G(0,k) > G(0,M)$. Let k' be the first such k. Then $G(0,k'\text{-}1) \leq G(0,M) < G(0,k')$. This then implies that $x_{k'}^*$, the point of intersection of $G(x,k'\text{-}1)$ and $G(0,k')$, is positive. This contradicts our earlier assertion that $x_k^* \leq 0$ for all $k > M$. Hence $G(0,k) \leq G(0,M)$ for all $k > M$. Since the slope of $G(x,k)$ is greater in absolute value than the that of $G(x,M)$, we conclude that $G(x,k) \leq G(x,M)$ for all $x \geq 0$. This proves the Proposition. ∎

Proof of Proposition 5.3. **(CASE III)**. In case III $\alpha\gamma < 1$, so the absolute values of the slopes of the $G(x,k)$ functions are decreasing in k. If for some k, $G(0,k) \geq G(0,k\text{-}1)$ then $x_k^* \leq 0$. Lemma 5.5(ii) would then imply that $x_k^* < 0$ for all large k. However, from (11) for all large k $G(0,k) \leq G(0,k\text{-}1)$, so $x_k^* \geq 0$ for all large k. This is a contradiction. Hence we conclude that $G(0,k) < G(0,k\text{-}1)$ for all $k = 1,2,\ldots$. This proves part (i) of the lemma.

Since $\alpha\gamma < 1$, $\alpha < 1/\gamma < 1$. From Lemma 5.5(i) the x_k^* 's are iterates of the b_2^{-1} function. This function has fixed point $\hat{x} = \sigma_\epsilon^2/(1\text{-}\alpha) > 0$, and slope $1/\alpha > 1$. From the definition of x_1^* it is easy to check that x_1^* can never be equal to $\hat{x}$. Iterates of this function from any $x_1^* \neq \hat{x}$ converge monotonically to either $+\infty$ or $-\infty$. Since we have just shown that the x_k^*'s are positive for all large k, we conclude that the monotonic convergence is to $+\infty$, from which part (ii) of the proposition follows. The remaining parts of the proposition then follow immediately. ∎

Proof of Proposition 5.4. (CASE IV). In CASE IV $\alpha\gamma < 1$ and $G(0,0) = \ell(0) < r(0,0) = G(0,1)$. Hence x_1^* must be negative. From Lemma 5.5(ii) this implies that x_k^* is negative for all k. However from (11), $G(0,k) \leq G(0,k\text{-}1)$ for all large k. Since the slopes of the functions $G(x,k)$ are decreasing in k this implies that x_k^* is positive for all large k. This is a contradiction. Hence case 4 is incompatible with our assumptions (and in particular assumption A.2, from which (11) was derived). ■

Proof of Lemma 5.6: (i) Since for all k, $g_k(\hat{x}) = h_1(h_2(\hat{x})) = h_1(\hat{x}) < \hat{x}$ and $g_k(0) > 0$, we conclude that $x_k^{**} \epsilon (0,\hat{x})$. However, for any x in $(0,\hat{x})$, $g_k(x) < g_{k+1}(x)$. It is easy to see that this implies that $x_k^{**} < x_{k+1}^{**}$. This proves the first part of the lemma.

(ii) We argued in the previous paragraph that $x_k^{**} \epsilon (0,\hat{x})$ for all k. Hence $h_2^k(x_k^{**}) < \hat{x}$ which implies that $x_k^{**} \equiv h_1(h_2^k(x_k^{**})) < h_1(\hat{x})$ for all k, so $\lim_{k\to\infty} x_k^{**} \leq h_1(\hat{x})$. Also, $x_k^{**} \equiv h_1(h_2^k(x_k^{**})) \geq h_1(h_2^k(0))$, and $\lim_{k\to\infty} h_2^k(0) = \hat{x}$, so $\lim_{k\to\infty} x_k^{**} \geq h_1(\hat{x})$. Combining our results shows that $\lim_{k\to\infty} x_k^{**} = h_1(\hat{x})$. ■

Proof of Proposition 5.9: If the initial posterior variance lies in $[x_1^*, \infty)$ the action NO SWITCH will be chosen. The posterior variance will in finite time enter the set $[0, x_1^*)$. From lemma 5.6 and the hypothesis of CASE IIA′, $x_k^{**} \leq x_M \leq x_1^*$ for all $k \leq M$. Hence once the posterior variance process enters the set $[0, x_1^*)$, it will stay there forever after. In the set $[0, x_1^*)$ the optimal jump size is greater than or equal to one. ■

Proof of Proposition 5.10. (Case IIB): Under CASE II, and lemma 5.6 above $x_1^* < x_1^{**} < x_2^{**} < x_3^{**} < \ldots$. Whenever the posterior variance is in the set $[x_1^*, \infty)$ the action NO SWITCH is chosen. This causes the posterior variance to decrease monotonically. After some finite date the posterior variance will fall below x_1^*. At this time a jump of size j equal one or larger will be chosen. This causes the posterior variance to move toward x_j^{**}. Since $x_j^{**} > x_1^*$ for all $j \geq 1$, eventually the posterior varaince will enter again the set $[x_1^*, \infty)$, and the process described above is repeated; ad infinitum. ■

Proof of Proposition 5.11. (CASE IIIA): (i) This should be obvious from fig. 9 of CASE III. (ii) If the initial posterior variance lies in the set (x_1^*, ∞), a jump size of $j \geq 1$ will be chosen. The hypothesis of CASE IIIA and lemma 5.6 imply that $x_1^* < x_k^{**}$ for all k. Hence if the initial posterior variance lies in the set (x_1^*, ∞) and $j \geq 1$ is chosen the posterior variance process will remain in the set (x_1^*, ∞). By induction we see that the posterior variance process will indeed be in the set (x_1^*, ∞) for each date and hence a jump size greater than or equal to one will be chosen at each date. ∎

Proof of Proposition 5.12. (CASE IIIB′): (i) Under the assumptions of CASE IIIB′ and lemma 5.6 $x_k \leq h_1(\hat{x}) < x_1^*$ for all k. Hence from any initial posterior variance, the posterior variance will enter the set $[0, x_1^*]$ in finite time. (ii) It should be obvious from fig. 9 that if the initial posterior variance lies in $[0, x_1^*]$ the optimal action is NO SWITCH (or jump size of zero) at each date. ∎

Proof of Proposition 6.2.: The continuity of $L(x)$, $R(x)$ and $V(x)$ follow from the dynamic programming arguments mentioned earlier and the continuity of $h(x)$, $h_1(x)$, $\ell(x)$ and $r(x)$.

To prove the other parts of the proposition we adapt the following arguments from Nyarko (1994). Let a <u>*-policy</u> to be a policy which is <u>independent</u> of both the date one posterior variance x and the history of observations at date t; a *-policy may however depend upon the date. (In particular,a *-policy is a deterministic rule for choosing actions as a function of the date.) Recall that Ψ is the set of all policies. Let Ψ^* be the set of all *-policies. Since $V(x) = \mathrm{Sup}_{\pi\epsilon\Psi} V_\pi(x)$, and since $\Psi^* \subseteq \Psi$,

$$V(x) \geq \mathrm{Sup}_{\pi\epsilon\Psi^*} V_\pi(x). \tag{14}$$

We showed in proposition 6.1 that the optimal <u>decision</u> at each date may be chosen to be a deterministic function of the status quo vintage and posterior variance of that vintage at each date. But the posterior variance over any vintage at any date is a deterministic function of the initial posterior variance x, and the sequence of decisions chosen up to that date. Hence the optimal decision at any date is a deterministic function of the initial date one variance, x, over vintage 1. Hence if we fix a posterior variance x over initial vintage 1, there will exist a *-policy which attains the same sum of discounted returns as the optimal policy. Hence for each x, $V(x) \leq \mathrm{Sup}_{\pi\epsilon\Psi^*} V_\pi(x)$. Combining this with (14) implies that

$$V(x) = \mathrm{Sup}_{\pi\epsilon\Psi^*} V_\pi(x). \tag{15}$$

The use of any *-policy, π, results in a fixed deterministic sequence of decisions SWITCH or NO SWITCH over time and hence a fixed deterministic sequence of vintages chosen at each date. This sequence is independent of the date one posterior variance x. The sequence of variances of the chosen vintages at each date is a sequence of iterates of the maps $h(\cdot)$ and $h_1(\cdot)$ defined earlier. Now, both $h(x)$ and $h_1(x)$ are strictly increasing and strictly concave functions of x. Hence any finite sequence of iterates of h and h_1 are also strictly increasing and strictly concave in x. Hence, the same is true of the variance of the vintage, $m(t)$, used at any given date t under the *-policy π, $\mathrm{Var}_t\, \theta_{m(t)}$. Now, the payoff at each date t is $\gamma^{m(t)}[1 - \mathrm{Var}_t\, \theta_{m(t)} - \sigma_w^2]$. The discounted sum of payoffs, $V_\pi(x)$ is a strictly convex and strictly decreasing function of x. Eq. (15) therefore implies that $V(x)$ is the supremum of a collection of strictly convex and strictly decreasing functions of x. It is easy to verify that since the supremum in (15) is attained for each x, $V(x)$ is therefore also a strictly convex and strictly decreasing function of x.

Next define Ψ^{*L} (resp. Ψ^{*R}) to be the set of all *-policies which choose the action NO SWITCH (resp. SWITCH) at date 1. Following the previous arguments it is easy to show that

$$L(x) = \mathrm{Sup}\,\{V_\pi(x)\colon \pi \,\epsilon\, \Psi^{*L}\} \quad \text{and} \quad R(x) = \mathrm{Sup}\,\{V_\pi(x)\colon \pi \,\epsilon\, \Psi^{*R}\},$$

and so $L(x)$ and $R(x)$ are each strictly decreasing and strictly convex in x. ∎

Proof of Proposition 6.3.: Let x and δ be as in the proposition. Let π denote the policy of choosing the action NO SWITCH in each period, and let $V_\pi(x)$ denote the return from this policy. Let π' denote the policy where the action SWITCH is chosen at date 1, but where in each subsequent period the action NO SWITCH is chosen. Let $V_{\pi'}(x)$ denote the return from this policy. Clearly it suffices to show that $V_\pi < V_{\pi'}$.

Under the policy π, the date t *prior* variance over the date t status quo technology (which of course is vintage 1) is $h_1^{t-1}(x)$ where $h_1^0(x) \equiv x$ and $h_1^{t-1}(x)$ is the t-1 times iterate of the function h_1 from x. Under the policy π', the date 1 prior over vintage 2 is $\alpha x + \sigma_\epsilon^2$. The date 2 prior variance over vintage 2 is $h_1(\alpha x + \sigma_\epsilon^2)$; and the date t prior variance over vintage 2 is $h_1^{t-1}(\alpha x + \sigma_\epsilon^2)$. Hence the difference in discounted payoffs of the two policies is

$$V_{\pi'}(x) - V_\pi(x) = \sum_{t=1}^{\infty} \delta^{t-1} \Delta_t \quad \text{where}$$

$$\Delta_t \equiv \gamma[1-\sigma_w^2 - h_1^{t-1}(\alpha x + \sigma_\epsilon^2)] - [1-\sigma_w^2 - h_1^{t-1}(x)] = (\gamma-1)(1-\sigma_w^2) + [h_1^{t-1}(x) - \gamma h_1^{t-1}(\alpha x + \sigma_\epsilon^2)]$$

Note that $r(x) - \ell(x) = \Delta_1$, so $\Delta_1 \geq 0$ by assumption. We will now show that $[h_1^{t-1}(x) - \gamma h_1^{t-1}(\alpha x + \sigma_\epsilon^2)]$ (which is negative) is strictly monotone increasing in t (so is becoming less negative as t increases). This will imply that $\Delta_t > 0$ for all $t > 1$. This in turn will prove that $V_{\pi'}(x) - V_\pi(x) > 0$ which will prove this proposition. Now, we may write $[h_1^{t-1}(x) - \gamma h_1^{t-1}(\alpha x + \sigma_\epsilon^2)] = A_t + B_t$ where $A_t \equiv [h_1^{t-1}(x) - h_1^{t-1}(\alpha x + \sigma_\epsilon^2)]$ and $B_t \equiv -(\gamma-1) h_1^{t-1}(\alpha x + \sigma_\epsilon^2)$. Since h_1 has slope everywhere less than one, $|A_t|$ is non-increasing in t. Since for $x \leq \hat{x}$, $h_1^{t-1}(x) < h_1^{t-1}(\alpha x + \sigma_\epsilon^2)$ we conclude that A_t is non-decreasing in t. Since $h_1^{t-1}(\alpha x + \sigma_\epsilon^2)$ is strictly decreasing in t, B_t is strictly increasing in t. Hence $A_t + B_t$ is strictly increasing in t. ∎

Proof of corollary 6.4: If from $x=0$ the action NO SWITCH is chosen then the posterior variance of the status quo technology will be $x=0$ in the next period as well. Hence from the stationarity of the problem, if from $x=0$ it is optimal to choose the action NO SWITCH at date 1, then it is optimal to choose that action in each and every subsequent date. However Proposition 6.3 implies that it is ***not*** optimal to choose the action NO SWITCH in each period. Hence the optimal action from $x=0$ is SWITCH. ∎

Proof of Proposition 6.5: The variance from period 2 onwards always lies between zero and σ_w^2. Hence the discounted future sum of returns (from dates 2 and onwards) is bounded between $V(0)$ and $V(\sigma_w^2)$. However, as $x \to \infty$ the difference between the current period rewards, $r(x)$ and $\ell(x)$ goes to infinity. Hence, there exists an $\bar{x}$ such that for all $x > \bar{x}$, $| r(x)-\ell(x) | > \mathrm{Sup}_x | \delta V(h_1(x))-V(h(x)) |$. Using the Bellman equation therefore proves the proposition. ∎

Proof of Proposition 6.6: Recall that a *-policy is nothing other than a sequence of decisions, SWITCH or NO SWITCH, for each date, with the decisions being made as a function of the date but independently of the posterior variance at that date. We begin with the following claim:

Claim: Define $\bar{M} \equiv 1/[1 - \mathrm{Max}\{\alpha, 1/\gamma\}]$. Fix any two initial date 1 posterior variances x' and x''. **i.** For any fixed *-policy π, $| V_\pi(x') - V_\pi(x'') | \leq \bar{M} | x' - x'' |$; and
ii. $V(x')-V(x'') \geq -\bar{M} | x'-x'' |$.

Proof of Claim: (i) Let π, x' and x'' be as in the claim. Let $m(t)$ denote the vintage of the technology chosen at date t and let $\sigma_t'^2$ (resp. $\sigma_t''^2$) be the prior variance at the beginning of date t of the vintage chosen at date t under the *-policy π from initial date 1 posterior variance x' (resp. x''). Define for each t,

$$w_t \equiv \delta^{t-1}\gamma^{m(t)-1} | (1 - \sigma_t'^2 - \sigma_w^2) - (1 - \sigma_t''^2 - \sigma_w^2) | = \delta^{t-1}\gamma^{m(t)-1} | \sigma_t'^2 - \sigma_t''^2 | .$$

Suppose that at date t the action SWITCH is chosen. Then $m(t+1) = m(t) + 1$. Also,

$\sigma_{t+1}'^2 = \alpha h_1(\sigma_t'^2) + \sigma_\epsilon^2$ and $\sigma_{t+1}''^2 = \alpha h_1(\sigma_t''^2) + \sigma_\epsilon^2$, so $|\sigma_{t+1}'^2 - \sigma_{t+1}''^2| \leq \alpha |\sigma_t'^2 - \sigma_t''^2|$. Hence

$$w_{t+1} = \delta^t \gamma^{m(t+1)-1} |\sigma_{t+1}'^2 - \sigma_{t+1}''^2| = \alpha(\delta\gamma)\delta^{t-1}\gamma^{m(t)-1} |\sigma_t'^2 - \sigma_t''^2| \leq \text{Max}\{\alpha, 1/\gamma\}.w_t,$$

where we use the fact that $\delta\gamma < 1$ for the last inequality above. Alternatively, suppose that at date t the action NO SWITCH is chosen. Then $m(t+1) = m(t)$. Also, $\sigma_{t+1}'^2 = h_1(\sigma_t'^2)$ and $\sigma_{t+1}''^2 = h_1(\sigma_t''^2)$ so (since $\partial h_1(x)/\partial x < 1$), $|\sigma_{t+1}'^2 - \sigma_{t+1}''^2| \leq |\sigma_t'^2 - \sigma_t''^2|$. Hence

$$w_{t+1} \equiv \delta^t \gamma^{m(t+1)-1} |\sigma_{t+1}'^2 - \sigma_{t+1}''^2| = \delta.\delta^{t-1}\gamma^{m(t)-1} |\sigma_t'^2 - \sigma_t''^2| \leq \text{Max}\{\alpha, 1/\gamma\}.w_t$$

where again we use the fact that $\delta < 1/\gamma$ for the last inequality. We therefore see that regardless of the decision chosen at date t, $w_{t+1} \leq \text{Max}\{\alpha, 1/\gamma\}.w_t$. Hence, $w_t \leq (\text{Max}\{\alpha, 1/\gamma\})^t w_1$ for all $t > 1$, so

$$\sum_{t=1}^{\infty} w_t \leq w_1/[1 - \text{Max}\{\alpha, 1/\gamma\}] = w_1\bar{M}.$$

Now $w_1 = \gamma^{m(1)-1} |\sigma_1'^2 - \sigma_1''^2|$. If the decision NO SWITCH was chosen at date 1 then $m(1) = 1$, $\sigma_1'^2 = x'$ and $\sigma_1''^2 = x''$ so $w_1 = |x' - x''|$. If the decision SWITCH was chosen at date 1 then $m(1) = 2$, $\sigma_1'^2 = \alpha x' + \sigma_\epsilon^2$ and $\sigma_1''^2 = \alpha x'' + \sigma_\epsilon^2$ so $w_1 \leq \alpha |x' - x''| \leq |x' - x''|$. In either case we see that $w_1 \leq |x' - x''|$. Putting this in the above equation and recalling the definition of w_t implies the conclusion of part (i) the claim.

(ii) Let π be the *-policy which is optimal from initial posterior x''. Then $V(x') \geq V_\pi(x')$ and $V(x'') = V_\pi(x'')$. Hence $V(x')-V(x'') \geq V_\pi(x')-V_\pi(x'')$. An application of part (i) of this lemma therefore proves part (ii). ∎

Proof of Proposition 6.6 (cont'd): Since $x \geq 0$, $h(x) \leq \sigma_w^2$ and $h_1(x) \geq 0$, and since V(.) is monotone decreasing,

$$V(h(x)) - V(h_1(x)) \geq V(\sigma_w^2) - V(0). \tag{16}$$

Part (ii) of the claim above implies that $V(\sigma_w^2) - V(0) \geq -\bar{M}\sigma_w^2$. (16) therefore implies that for all $x \geq 0$,

$$V(h(x)) - V(h_1(x)) \geq -\bar{M}\sigma_w^2. \tag{17}$$

Fix any $\bar{\sigma}^2 > \sigma_w^2$. Define

$$M = \text{Max}\ \{\ |\ \ell(x) - r(x)\ |\ :\ x \in [0, \bar{\sigma}^2]\}. \tag{18}$$

Then $M < \infty$.

Suppose that $r(x^{**}) > 0$. Then there exists a $\xi > 0$ sufficiently small such that $r(x^{**} + \xi) > \xi$. We now compute the return to the policy, π_{SWITCH}, of switching in each period from initial posterior $x = \sigma_w^2$. Let $\{x_t\}_{t=1}^{\infty}$ be the associated posterior variance process from π_{SWITCH}. Since $x_t \to x^{**}$, there exists a $T < \infty$ such that for all $t > T$, $x_t < x^{**} + \xi$, so $r(x_t) > \xi$. If M is as in (18) above, the discounted sum of returns to this policy will then necessarily exceed $-TM + \sum_{t=T+1}^{\infty}(\delta\gamma)^{t-1}\xi = -TM + \xi(\delta\gamma)^T/(1 - \delta\gamma)$. This tends to infinity as $\delta \to 1/\gamma$. Hence $V(\sigma_w^2)$ tends to infinity as $\delta \to 1/\gamma$.

Since for all $x \geq 0$, $h(x) \leq \sigma_w^2$, $V(h(x)) \geq V(\sigma_w^2)$ so $V(h(x))$ tends to ∞ as $\delta \to 1/\gamma$. Hence we may choose a $\bar{\delta} < 1/\gamma$ such that for all $\delta \geq \bar{\delta}$ and for all $x \geq 0$,

$$\delta(\gamma - 1)V_\delta(h(x)) > \delta\bar{M}\sigma_w^2 + M, \tag{19}$$

where we subscript the value function V by δ to emphasize its dependence on δ, and where $\bar{M}$ and M are as in (17) and (18), respectively, (and are independent of δ). Using (17) and

(18) in (19) implies that for all such δ, and for all $x \in [0, \bar{\sigma}^2]$,

$$\delta\gamma V_\delta(h(x)) - \delta V_\delta(h_1(x)) = \delta(\gamma - 1)V_\delta(h(x)) + \delta[V_\delta(h(x)) - V_\delta(h_1(x))] > M + \delta\bar{M}\sigma_w^2 - \delta\bar{M}\sigma_w^2$$
$$= M > \ell(x) - r(x),$$

so

$$R(x) = r(x) + \delta\gamma V_\delta(h(x)) > \ell(x) + \delta V_\delta(h_1(x)) = L(x).$$

Hence the optimal action from such δ and x is SWITCH. ▮

Proof of Proposition 6.8: It suffices to show that from any initial posterior variance the agent chooses each action, SWITCH and NO SWITCH, infinitely often.

Suppose that from initial posterior variance x^{**} the optimal action is SWITCH. Then since the posterior variance in the next period will remain at x^{**}, it will be optimal to SWITCH in each and every period. Since $r(x^{**}) < 0$, this will result in a negative payoff in each period. As δ converges to $1/\gamma$ this payoff can be easily seen to converge to $-\infty$. If on the other hand the agent chooses the action NO SWITCH in each period, the agent will have a return of at least $\ell(x^{**})/(1-\delta)$ which converges to the finite number $\ell(x^{**})/(1-1/\gamma)$ as $\delta \to 1/\gamma$. Hence there will exist a $\bar{\delta}$ in $(0, 1/\gamma)$ such that for all $\delta \in (\bar{\delta}, 1/\gamma)$ it is not optimal to SWITCH from x^{**}. Under obvious continuity of the $L(x)$ and $R(x)$ functions the optimal action within a neighborhood of x^{**} is also NO SWITCH. If the agent chooses the action SWITCH for all but finitely many periods, the variance process will converge to x^{**}. This will contradict the earlier assertion that from a neighborhood of x^{**} the optimal action is NO SWITCH. Hence the agent will choose the action NO SWITCH infinitely often.

Using Corollary 6.4, from initial posterior variance $x = 0$, and hence from a neighborhood of $x = 0$, the optimal action is SWITCH. If from some initial posterior variance $x_1 > 0$ the agent chooses the action NO SWITCH for all but finitely many dates then the posterior variance process will converge to zero. This will contradict the fact that from a neighborhood of zero the optimal action is SWITCH. Hence the action SWITCH will be chosen infinitely often.

In summary we have shown that from any initial posterior variance the optimal policy will choose both actions, SWITCH and NO SWITCH, infinitely often. Hence we have "cycles." ▮

Proof of Proposition 6.9. (small δ): We prove part (ii) first. Fix any $\xi > 0$. Define

$$\kappa(x) \equiv | \ell(x) - r(x) | \quad \text{for all } x \geq 0.$$

Under the assumption that the parameter values lie in the generic set G both $\ell(x)$ and $r(x)$ are linear functions with different slopes and different intercepts and with a strictly positive intersection x^*. It is therefore easy to see that

$$\kappa(x) \geq \kappa(x^*+\xi) = \kappa(x^* - \xi) > 0 \text{ for all } x \in [0, x^* - \xi) \cup (x^* + \xi, \infty). \tag{20}$$

Now, $h_1(x) \equiv \sigma_w^2 x/(x + \sigma_w^2) \leq \sigma_w^2$ and hence $h(x) \equiv h_1(h_2(x)) \leq \sigma_w^2$. Hence regardless of the current period decision, SWITCH or NO SWITCH, the next period posterior variance will be bounded above by σ_w^2. This will be true for the posterior variance at each date. Hence absolute value of the expected payoff in each period is uniformly bounded, by some number K say. Hence $| V_\delta(h_1(x)) | \leq K/(1 - \delta)$ and $| \gamma V_\delta(h(x)) | \leq \gamma K/(1 - \delta)$, where we now index the value function by the subscript δ to denote its dependence on δ. This in turn implies that $\mathrm{Sup}_{x\geq 0}\, \delta\, | V_\delta(h_1(x)) - \gamma V_\delta(h(x)) | \to 0$ as $\delta \to 0$. In particular we may choose a $\bar{\delta} > 0$ such that

$$\delta\, | V_\delta(h_1(x)) - \gamma V_\delta(h(x)) | < \kappa(x^* - \xi)/2 \quad \text{for all } \delta < \bar{\delta}. \tag{21}$$

Combining this with (20) above implies that

$|\ \ell(x) - r(x)\ | \;>\; \delta\ |\ V_\delta(h_1(x)) - \gamma V_\delta(h(x))\ | \quad \forall\delta < \bar{\delta}$ and $\forall x$ in $[0, x^* - \xi) \cup (x^* + \xi, \infty)$.

It is then easily verified that this in turn means that

$$\begin{aligned} &\ell(x) > r(x) \text{ implies that } \ell(x) + \delta V_\delta(h_1(x)) > r(x) + \gamma V_\delta(h(x)) \text{ and} \\ &\ell(x) < r(x) \text{ implies that } \ell(x) + \delta V_\delta(h_1(x)) < r(x) + \gamma V_\delta(h(x)). \end{aligned} \tag{22}$$

Note that $L(x) \equiv \ell(x) + \delta V_\delta(h_1(x))$ (resp. $R(x) \equiv r(x) + \gamma V_\delta(h(x))$) is the return to the agent that chooses the action NO SWITCH (resp. SWITCH) in the current period and behaves optimally in each subsequent period. (22) therefore implies that the decision which maximizes the myopic return also maximizes the infinite horizon return with discount factor δ. This of course is part (ii) of the proposition.

Part (i): To prove part (i) of the proposition observe that $\kappa(x) \equiv |\ \ell(x) - r(x)\ |$ is monotone non-decreasing in x in CASES I and IV. Further, $[1 - \sigma_w^2] \neq \gamma[1 - \sigma_\epsilon^2 - \sigma_w^2]$, $\kappa(0) \equiv |\ \ell(0) - r(0)\ | > 0$. Hence we may replace (20) with $\kappa(x) \equiv |\ \ell(x) - r(x)\ | \geq \kappa(0) > 0$, and proceed analogously to the proof of part (ii) above to obtain the proof of part (i) of the proposition. ∎

Proof of Lemma 7.1.: In the myopic problem from initial posterior variance $x=0$ the agent will never choose a jump size greater than $K+1$ since this will result in a negative payoff while choosing no jump in each period results in a payoff of $1-\sigma_w^2>0$ in each period. The return to choosing a jump size of $k \leq K+1$ results in a payoff no greater than $\gamma^{K+1}[1-\sigma_w^2]$. From the definition of x_ψ it is easy to see that $\gamma^{K+1}[1-\sigma_w^2] \leq x_\psi$. Hence $V^1(0) \leq x_\psi$.

We proceed by induction: suppose that for some $T=1,2,\ldots,$ we have shown that $V^T(0) \leq x_\psi$. The value function of the $T+1$ horizon problem from initial posterior variance $x=0$ is:

$$V^{T+1}(0) = \mathrm{Max}_{k=0,1,2,\ldots}\ \gamma^k\{1-\sigma_w^2-h_2^k(0)+\delta V^T(h_1(h_2^k(0)))\}.$$

If a jump size of $k>K+1$ is optimal for this $T+1$ horizon problem from $x=0$, then the return will be no greater than $\gamma^k\{1-\sigma_w^2-h_2^k(0)+\delta V^T(0)\}$, which from the induction hypothesis is no greater than $\gamma^k\{1-\sigma_w^2-h_2^k(0)+\delta x_\Psi\}$. This latter expression is, from (21), negative for all $k>K+1$. However, if a jump size of zero is chosen at each date the return at each date will be $1-\sigma_w^2$ which is positive. Hence it is not optimal to choose a jump size $k>K+1$. With this result we may conclude from the definition of V^{T+1} above that

$$V^{T+1}(0) \leq \gamma^{K+1}(1-\sigma_w^2)+\delta\gamma^{K+1}V^T(0) \equiv \Psi(V^T(0)).$$

Since $V^T(0) \leq x_\Psi$, we conclude from the obvious properties of Ψ that $\Psi(V^T(0)) \leq x_\Psi$. So $V^{T+1}(0) \leq x_\Psi$. ▮

Appendix 2. Verification that all the CASES I-IV are non-empty.

It should be obvious that there exist parameter values such that CASES I and IV hold. Set $\sigma_\epsilon^2 = \sigma_w^2 = 0.25$. Then

$$[\gamma(1 - \sigma_\epsilon^2 - \sigma_w^2) - (1 - \sigma_w^2)] = [0.5\gamma - 0.75] = [2\gamma - 3]/4 \text{ and}$$

$$x^* = [\gamma(1 - \sigma_\epsilon^2 - \sigma_w^2) - (1 - \sigma_w^2)]/[\gamma\alpha - 1] = [2\gamma - 3]/[4(\gamma\alpha - 1)].$$

Also,

$$x^{**} = [\sigma_w^2(\alpha x^{**} + \sigma_\epsilon^2)]/[\sigma_w^2 + (\alpha x^{**} + \sigma_\epsilon^2)]$$

Hence $\sigma_w^2\sigma_\epsilon^2/(\sigma_w^2 + \sigma_\epsilon^2) \leq x^{**} \leq \sigma_w^2$ so

$$1/8 \leq x^{**} \leq 1/4. \tag{23}$$

CASE IIA $x^* > x^{**}$. Fix any α. Let γ tend to infinity. Then $x^* = 0.25[2\gamma - 3]/[\alpha\gamma - 1] \rightarrow 1/(2\alpha)$. Hence for any fixed $\alpha < 1$, we may choose a γ sufficiently large so that $(\alpha\gamma) > 1$, $0.25[2\gamma - 3] > 1$ and $x^* > 1/(2\alpha) > 1/4 > x^{**}$, so that we are in CASE IIA.

CASE IIB $x^* < x^{**}$. Take a sequence $\{\gamma_k\}^{\infty}_{k=1}$ converging to 3/2 from above. Let α_k be any sequence converging to 3/4 from above. Then $\alpha_k\gamma_k \rightarrow 9/8 > 1$ and $0.25(2\gamma_k - 3) > 0$ for all k. Hence we are in CASE II for all k sufficiently large. Further, $x_k^* = 0.25[2\gamma_k - 3]/[\alpha_k\gamma_k - 1] \rightarrow 0$ so $x_k^* < 1/8 < x^{**}$ for all k sufficiently large, which implies CASE IIB.

Case IIIA: $x^* < x^{**}$. Now consider a sequence $\{\gamma_k\}^{\infty}_{k=1}$ converging to 3/2 from below. Define $\alpha_k = 1/2\gamma_k$, in which case α_k converges to 1/3. Also $\alpha_k\gamma_k = 1/2 < 1$ and $2\gamma_k - 3$ converges to zero from below so for all k sufficiently large we are in CASE III. Now, in this case $x_k^* = 0.25[2\gamma_k - 3]/[\alpha_k\gamma_k - 1] \rightarrow 0$ so for k sufficiently large $x^* < 1/8 < x^{**}$ so we are in CASE IIIA.

Case IIIB: $x^* > x^{**}$. Fix any $\alpha < 1$. Pick any sequence of γ's converging to one from above. Then eventually, $(\alpha\gamma) < 1$ and $[\gamma(1 - \sigma_\epsilon^2 - \sigma_w^2) - (1 - \sigma_w^2)] = 0.25[2\gamma - 3] < 0$ so we are indeed in case III. Also $x^* = 0.25[2\gamma - 3]/(\alpha\gamma - 1)$ converges to $0.25/(1 - \alpha) > 0.25$, so from equation (23) we conclude that $x^* > x^{***}$ so we are in CASE IIIB.

Each of the above cases can coexist with the restriction, made in proposition 6.6, that $r(x^{**}) \geq 0$. The following six sets of parameters do the job:

Case I: $\sigma_w^2 = 0.30$, $\sigma_t^2 = 0.25$, $\alpha = 0.75$, $\gamma = 1.55$.
Case IIA: $\sigma_w^2 = 0.25$, $\sigma_t^2 = 0.25$, $\alpha = 0.75$, $\gamma = 1.55$.
Case IIB: $\sigma_w^2 = 0.25$, $\sigma_t^2 = 0.25$, $\alpha = 0.75$, $\gamma = 1.54$.
Case IIIA: $\sigma_w^2 = 0.20$, $\sigma_t^2 = 0.30$, $\alpha = 0.50$, $\gamma = 1.55$.
Case IIIB: $\sigma_w^2 = 0.20$, $\sigma_t^2 = 0.30$, $\alpha = 0.60$, $\gamma = 1.55$.

Case IV: $\sigma_w^2 = 0.10$, $\sigma_\ell^2 = 0.30$, $\alpha = 0.60$, $\gamma = 1.55$.

Appendix 3. Proof that the Conditions in Proposition 6.8 Rule Out Case IV. It is enough to show that $r(x^{**}) < 0$ and $\gamma(1 - \sigma_w^2 - \sigma_\epsilon^2) > (1 - \sigma_w^2)$ is impossible when $\alpha\gamma < 1$. First, note that

$$x^{**} = \frac{\sigma_w^2(\alpha x^{**} + \sigma_\epsilon^2)}{\sigma_w^2 + (\alpha x^{**} + \sigma_\epsilon^2)}$$

the solution of which is

$$x^{**} = \frac{\{ -(\sigma_w^2 + \sigma_\epsilon^2 - \alpha\sigma_w^2) + [(\sigma_w^2 + \sigma_\epsilon^2 - \sigma_w^2\alpha)^2 + 4\alpha\sigma_w^2\sigma_\epsilon^2]^{1/2} \}}{2\alpha} .$$

The negative root is inadmissable because $x \geq 0$. Now $r(x^{**}) = \gamma(1 - \sigma_w^2 - \sigma_\epsilon^2 - \alpha x^{**}) < 0$ implies that $1 - \sigma_w^2 - \sigma_\epsilon^2 < \alpha x^{**}$. Substitution for x^{**} into the above inequality and (very tedious) rearrangement gives

$$\begin{aligned} \gamma(1 - \sigma_w^2 - \sigma_\epsilon^2) &< \gamma\alpha\sigma_w^2(1 - \sigma_w^2) \\ &< \sigma_w^2(1 - \sigma_w^2) && \text{since } \gamma\alpha < 1 , \\ &\leq (1 - \sigma_w^2) && \text{since } \sigma_w^2 \leq 1 . \end{aligned}$$

(We assume $\sigma_w^2 \leq 1$ since otherwise net outputs are always negative.) So $r(0) < \ell(0)$. ∎

Appendix 4. The No-Recall Assumption:

The analysis of this paper assumed the following no-recall constraint: An agent who chooses a vintage n at date t, can not choose a lower vintage $n' < n$ at any future date $t' > t$. We now proceed to show that if we are in the myopic model, it is never optimal to exercise the recall option and choose a vintage that was previously passed over. Our conjecture is that the same result is true in the positive discount factor problem. Our conjecture is given support by the fact that we have a proof that it is not optimal to exercise the recall option in the two-horizon version of the problem with positive discount factors.

To proceed we will require some notation: Suppose that at date 1 the variance-covariance matrix over $(\theta_1,\theta_2,\ldots)$ is given by $\Sigma=(\{x_{ij}\}_{i,j=1,2,\ldots})$.

Lemma A4.0: Then after observation of $y_m=\theta_m+w_m$ the posterior variance of θ_k is given by

$$\mathrm{Var}\ (\theta_k \mid y_m) = x_{kk}-x_{km}^2/[x_{mm}+\sigma_w^2].$$

Proof: Let Σ^* denote the posterior variance-covariance matrix after observation of $y_m=\theta_m+w_m$. Then from Chow (1983, p. 13), $\Sigma^*=\Sigma-\Sigma_{12}\Sigma_{22}^{-1}\Sigma_{21}'$, where Σ_{12} is the matrix of covariances between y_m and $(\theta_1,\theta_2,\theta_3,\ldots)$ and where Σ_{22} is the variance of y_m. Now, $\mathrm{Var}\ y_m = x_{mm}+\sigma_w^2$; for any $j=1,2,\ldots$, $\mathrm{Cov}\ (\theta_j,y_m)=E[(\theta_j-E\theta_j)(y_m-Ey_m)]=E[(\theta_j-E\theta_j)(\theta_m+w_m-E\theta_m)]=x_{mj}$. Hence applying the formula we see that

$$\Sigma^*=\Sigma-(x_{1m},x_{2m},\ldots)(x_{1m},x_{2m},\ldots)'/[x_{mm}+\sigma_w^2].$$

The (k,k) element of this matrix can easily be seen to be $x_{kk}-x_{km}^2/(x_{mm}+\sigma_w^2)$. ∎

Let q_1 and q_2 denote the dates 1 and 2 output levels respectively. Let $Q_1(<k>)\equiv E[q_1 \mid <k>]$ denote the expected date 1 output level if vintage k is chosen at date 1. Let $Q_2(<k_1,k_2>)\equiv E[q_2 \mid <k_1,k_2>]$ denote the date 2 output level if vintage k_1 is chosen at date 1 and vintage k_2 is chosen at date 2. Let $R(<k_1,k_2>)\equiv Q_1(<k_1>)+\delta Q_2(<k_1,k_2>)$ be the discount sum of net outputs when vintage k_1 is chosen at date 1 and vintage k_2 is chosen at date 2.

Lemma A4.1: Suppose that it is optimal to choose vintage n at date 1 and vintage 1 at date 2.

Then $R(<n,1>)-R(<1,n>) \geq 0$.

Proof: Suppose agent A chooses vintage n followed by vintage 1, while agent B does the reverse and chooses vintage 1 followed by vintage n. If each began with the same date 1 prior each agent will have at the beginning of date 3 the same variance-covariance matrix since they receive the same information, although in different orders. In the no-jump model each agent will also have at the beginning of date 3 the same maximal vintage chosen. Hence from date 3 onwards the decision problems of the two agents will be identical. Suppose each agent chooses the optimal policy from dates 3 onwards. By hypothesis of the lemma, agent A's policy is optimal. Since agents A and B will have the same return from dates 3 onward, agent A's return in the first two periods can not be less than that of B's in the first two periods. The conclusion of the lemma follows from this observation. ▮

Lemma A4.2: $R(<n,1>)-R(<1,n>)=[1-\delta+\delta\zeta][Q(<n>)-Q(<1>)]-\delta\zeta(\lambda^{n-1}-1)$

where $\zeta=x_{1n}^2/[(x_{nn}+\sigma_w^2)(x_{11}+\sigma_w^2)]>0$.

Proof: Since

$\mathrm{Var}\,[\theta_n \mid <1>]=x_{nn} - x_{1n}^2/(x_{11}+\sigma_w^2)$ and

$\mathrm{Var}\,[\theta_1 \mid <n>]=x_{11} - x_{1n}^2/(x_{nn}+\sigma_w^2)$, one may easily verify that

$$R(<n,1>) = \lambda^{n-1}\{1-\sigma_w^2-x_{nn}\} + \delta\{1-\sigma_w^2-x_{11}+ x_{1n}^2/(x_{nn}+\sigma_w^2)\} \text{ and}$$
$$R(<1,n>)=\{1-\sigma_w^2-x_{11}\} + \delta\lambda^{n-1}\{1-\sigma_w^2-x_{nn}+ x_{1n}^2/(x_{nn}+\sigma_w^2)\}$$

from which the lemma follows immediately. ▮

Lemma A4.3: $R(<n,1>)-R(<1,n>)\geq 0$ implies that $Q_1(<n>)\geq Q_1(<1>)$.

Proof: This follows immediately from lemma A4.2. ▮

Lemma A4.4: $Q_1(<n>) \geq Q_1(<1>)$ implies that $Q_2(<n,n>) > Q_2(<n,1>)$.

Proof:

$Q_2(<n,n>)-Q_2(<n,1>)=\lambda^{n-1}\{1-\sigma_w^2 - x_{nn} + x_{nn}^2/(x_{nn}+\sigma_w^2)\} - \{1-\sigma_w^2 - x_{11} + x_{1n}^2/(x_{nn}+\sigma_w^2)\}$.

The Cauchy-Schwartz inequality implies that $x_{12}^2 \leq x_{nn}x_{11}$. Hence the above is

$$\geq (\lambda^{n-1}-1)(1-\sigma_w^2) - (\lambda^{n-1}x_{nn}-x_{11})(1- x_{nn}/(x_{nn}+\sigma_w^2)) \qquad (*)$$

$$= Q_1(<n>)-Q_1(<1>) + (\lambda^{n-1}x_{nn}-x_{11})x_{nn}/(x_{nn}+\sigma_w^2). \qquad (**)$$

If $(\lambda^{n-1}x_{nn}-x_{11}) \leq 0$ then the RHS of (*) is non-negative from which the lemma follows. If $(\lambda^{n-1}x_{nn}-x_{11}) > 0$ then the lemma follows from (**). ∎

Proposition A4.5 (No Recall): Suppose that we are either in the (a) myopic model or (b) two-horizon model. Then an agent will never use the recall option.

Proof: If the agent uses the recall option then (by rescalling) we may assume that at some date T the agent uses a vintage $n > 1$, and at date $T+1$ uses vintage 1.

(a) Suppose we are in the myopic model. Then at date T the return from using vintage n, $Q_T(<n>)$, must exceed that of using vintage 1, $Q_T(<1>)$. From lemma A4.4 this implies that the date $T+1$ return from using vintage n, $Q_{T+1}(<n,n>)$, exceeds that of using vintage 1, $Q_{T+1}(<n,1>)$. However, for the myopic model, this implies that vintage n is used in period $T+1$, which contradicts our earlier assertion.

(b) Now suppose that we are in the two period model. Then $R(<n,1>) \geq R(<1,n>)$. Lemmas A4.3 and A4.4 imply that $Q_2(<n,n>) > Q_2(<n,1>)$. But this in turn implies that for the two period model, at date $T+1$ vintage n is chosen. This contradicts our earlier assertion. ∎

References:

Chari, V.V., and Hugo Hopenhayn, "Vintage Human Capital," *Journal of Political Economy* 99, no. 6 (December 1991): 1142-1165

Chow, Gregory, *Econometrics*, New York, N.Y.: McGraw Hill (1983).

Lucas, Robert E., "Making a Miracle," *Econometrica* 61, no. 2 (March 1993): 251-72.

Nyarko, Yaw, "On the Convexity of the Value Function in Bayesian Optimal Control Problems," *Economic Theory*, (1994).

Parente, Stephen, "Technology Adoption, Learning by Doing, and Economic Growth," unpublished paper, Northeastern University, October 1991.

Prescott, Edward, "The Multi-Period Control Problem Under Uncertainty," *Econometrica* 40, no. 6 (November 1972): 1043-58.

Stokey, Nancy, "Human Capital, Product Quality, and Growth," *Quarterly Journal of Economics* 106 (1991): 587-616.

Wilson, Robert, "Informational Economies of Scale," *Bell Journal of Economics* (Spring 1975): 184-95.

Young, Alwyn, "Invention and Bounded Learning by Doing," *Journal of Political Economy* 101, no. 3 (June 1993): 443-72.