REVIEW OF ECONOMIC DYNAMICS **1,** 531–550 (1998)
ARTICLE NO. RD980014

Learning Curves and the Cyclical Behavior of Manufacturing Industries

Peter J. Klenow*

Graduate School of Business, University of Chicago, Chicago, Illinois 60637

Received June 1, 1997

Building on evidence that (a) productivity growth from learning by doing diminishes as experience accumulates with a technology and (b) learning by doing is largely specific to each production technology, this paper models a firm's decision of when to update its technology. The model implies that technology updates endogenously bring large drops in productivity. The model also implies that technology updates are more likely in a boom than in a recession since a high rate of production enables the firm to learn more quickly about the new technology. The forces in this model may help explain some features of plant and industry level data, such as the procyclicality of investment (including plant investment spikes) and the modest correlation between labor input and productivity. *Journal of Economic Literature* Classification Numbers: O31, L6, D92.

1. INTRODUCTION

A number of empirical studies of plant productivity indicate that when firms change their production technology, their productivity initially falls and then gradually rises to eventually overtake the level achieved with the old technology.[1] The drop in productivity suggests that production knowledge does not apply equally across the old and new production technologies, and the gradual rise in productivity suggests learning by doing.[2] If

* I am grateful to Boyan Jovanovic and two anonymous referees for many helpful comments.

[1] See the evidence presented in Cochran [14], Garg and Milliman [21], Baloff [5–7], Russell [34], and Pegels [32]. See also Yorukoglu [36] for evidence of the same pattern for firms investing in information technology.

[2] For surveys of empirical studies of learning by doing see Yelle [35] and Argotte and Epple [1]. Jovanovic and Nyarko [30] also discuss a number of studies. Bahk and Gort [3] find evidence of learning by doing in U.S. manufacturing plants. Examining firm-level data on individual semiconductor products, Irwin and Klenow [28] find evidence for learning by doing and its limited applicability across generations of chips.

1094-2025/98 $25.00

what firms learn is specific to a given production technology the question arises as to why firms update their technology at all. The answer could be that productivity gains from learning eventually slow, while the new technologies continue to get better and better. Eventually it is worth abandoning an old technology for a newer technology despite the lack of experience with the newer technology.[3]

Building on evidence that learning is specific to each production technology and yields substantial but diminishing productivity gains, this study models a firm's decision of when to update its process technology.[4] I emphasize three implications of the model. First, firm productivity endogenously falls when the firm updates its technology. The firm does not wait until operating the leading-edge technology with no experience matches the productivity of the technology currently used since learning is initially rapid with the new technology. Second, firms are more likely to update technology when demand is high than when demand is low. This is because a sustained high rate of production enables the firm to learn more quickly about the new technology. Third, the correlation between firm productivity and labor input is low because the rate of learning is high when productivity is relatively low and low when productivity is relatively high. The firm keeps labor input high as it commences production with a new technology because, although the current marginal product is low, there is much to learn.

The remainder of the paper proceeds as follows. In Section 2 I describe the model in detail. In Section 3 I then cast the firm's profit maximization problem as a dynamic program. In Section 4 I calibrate the model and solve it for optimal firm behavior in the face of demand shocks: when to update technology and how much labor to hire each period. I characterize the model's implications in Section 5, and compare them to features of manufacturing industry and plant level data in Section 6. The industry data comes from the NBER Manufacturing Productivity Database maintained

[3] A slowing rate of productivity gains from learning does not require that learning be literally bounded. Consider the typical formulation of productivity as a function of cumulative experience raised to a power less than 1. In this formulation learning is unbounded, but the rate of productivity gains from learning falls as experience accumulates. Asher [2], Conway and Schultz [15], and Baloff [5, 8] provide evidence that productivity gains from learning slow down as experience accumulates. Young [37] incorporates bounded learning in a model of endogenous growth.

[4] The environment explored by Zeckhauser [38] is quite close to the one examined here. Another related paper is Parente [31]. Parente also models technology adoption in the presence of firm-specific and technology-specific learning by doing (although not with variable labor input and demand shocks). His emphasis is on how capital markets facilitate technology adoption given consumers' desire to smooth consumption in the face of firm fluctuations. In contrast, the emphasis here is on calibrating a model to explore its potential usefulness in explaining industry cyclical behavior and idiosyncratic plant behavior.

by Bartelsman and Gray [10]. The plant level data is described in other studies, such as Baily, Hulten, and Campbell [4]. I offer conclusions in Section 7.

2. A MODEL WITH ENDOGENOUS TECHNOLOGY UPDATING

Consider a representative consumer maximizing utility over hours worked and over the consumption of a fixed measure M of goods, with momentary utility of the form

$$\mathrm{u}(C_t) = C_t^{1/(1-1/\epsilon)} - \eta N_t$$

with

$$C_t = \int_0^M Y_t(i)^{1-1/\epsilon}\, di.$$

Here $Y(i)$ is output of good $i \in [0, M]$ and ϵ is the elasticity of substitution between any pair of goods. Under these preferences the consumer's demand curve for good i is

$$P_t(i) = \left[\frac{Y_t(i)}{C_t}\right]^{-1/\epsilon}, \tag{2.1}$$

where $P_t(i)$ is the price of good i relative to that of the utility-maximizing composite C_t. The price elasticity of demand, conditional on C, equals the elasticity of substitution ϵ between goods in utility. The demand for good i is affected by the prices of other goods solely through the composite C. Conditional on a one-to-one mapping between goods and firms, firms affect each other only through their impact on C. Given a positive measure of firms, each firm views C as an exogenous demand process.

Each firm maximizes the present value of its cash flows

$$E_0 \sum_{t=0}^{\infty} (1+r)^{-t} [P_t Y_t - w_t N_t].$$

Here r denotes the real interest rate, w_t the wage and N_t denotes the firm's labor input. The firm hires labor in a competitive market, with the real wage growing at the rate μ_e of economywide average productivity growth. Elastic demand ($\epsilon > 1$) is, of course, necessary for the firm's problem to be well defined. "Demand" C is assumed to be the product of a stationary stochastic component and a nonstochastic trend component:

$$C_t = D_t T_t,$$
$$D_t = D_{t+1}^{\rho} e^{\omega_t}, \tag{2.2}$$
$$T_t = T_{t-1} e^{\mu_e},$$

where $|\rho| < 1$ and $\omega_t \sim$ i.i.d. $N(0, \sigma_\omega^2)$. The firm's production technology is

$$Y_t = Q_t^{\lambda} N_t, \tag{2.3}$$

where Q_t^{λ} is the level of productivity and Q_t is the firm's level of ''experience'' with the technology it is currently using. The firm's technology is τ periods old, so it was leading-edge τ periods ago. Each period the firm decides whether to continue with the old technology ($\tau_t = \tau_{t-1} + 1$) or adopt the leading-edge technology ($\tau_t = 0$).[5] The firm faces the following transition process for experience Q_t:

$$Q_t = \begin{cases} A_t^{1/\lambda}, & \text{if } \tau_t = 0, \\ \min\left(\gamma A_{t-\tau_t}^{1/\lambda}, e^{\mu/\lambda} Q_{t-1} + \dfrac{Y_{t-1}}{A_t^{\epsilon - 1/\lambda} T_t^{1-\epsilon}}\right), & \text{if } \tau_t = \tau_{t-1} + 1, \end{cases} \tag{2.4}$$

with

$$A_t = A_{t-1} e^{\mu}, \qquad \gamma > 1, \qquad 0 < \lambda < 1.$$

A_t represents the level of frontier technology in the industry, which advances at the rate μ per year. As the first line of (2.4) shows, when the firm adopts a leading-edge technology ($\tau = 0$) it does not operate the technology at maximum efficiency right away. The firm has $(\gamma - 1)$% to learn about operating the technology at efficiency γA_t. Each technology confronts the firm with a separate learning curve. As shown in the second term in the min expression of (2.4), with each technology firms learn by

[5] The firm never chooses $0 < \tau_t < \tau_{t-1} + 1$ because adopting an old technology is inferior to adopting the frontier technology. The firm either sticks with last period's technology or switches to the most advanced technology. This result depends critically on the assumption of no transferability of learning to ''nearby'' technologies. Lack of transferability is consistent with the evidence cited earlier for productivity drops, as well as with the evidence on semiconductor chips in Irwin and Klenow [28].

doing as their production Y feeds into higher experience Q.[6] The parameter λ determines the speed of learning: the higher is λ, the faster the firm closes the gap between its attained productivity and the maximum attainable productivity with the vintage. As shown in the first term in the min expression of (2.4), the learning is bounded above for each technology by $\gamma A_{t-\tau}$; there is only so much to learn. The more production experience the firm has with a technology, the less it has left to learn. The learning is specific to each technology, as the starting productivity with each technology is the same regardless of the level of experience with the old technology. Why do firms update their technologies at all in this environment? Because an old technology, even operated at peak efficiency, eventually becomes less efficient than the ever-advancing frontier technology, even operated with no experience.

3. SOLVING THE MODEL

The leading-edge technology A_t faced by the firm advances at the constant rate μ. Since Q starts at $A^{1/\lambda}$ with each technology, productivity Q^{λ} trends upward at the average rate of μ. Economywide average productivity growth occurs at the rate μ_e, driving economywide real wage growth at the same rate. Since the firm faces real wages rising at rate μ_e, firm labor input N grows (on average) at the rate $(\epsilon - 1)(\mu - \mu_e)$.[7] A firm with faster productivity growth than the economy overall will draw in labor, more so the more price-elastic is demand. Combining (2.3) and growth in the leading-edge technology at rate μ, the firm's real output will grow at the rate $\mu + (\epsilon - 1)(\mu - \mu_e)$. Equation (2.1) then implies that the firm's relative price will grow at the rate $(\mu_e - \mu)$. An industry with rapidly advancing technology $(\mu > \mu_e)$, such as personal computers, will see its relative price fall over time. Finally, the firm's "nominal" output will grow at the rate $\mu_e + (\epsilon - 1)(\mu - \mu_e)$.

Using these average growth rates, the firm's problem can be reexpressed in terms of stationary variables. Defining $\hat{Q}_t = Q_t/A_t^{1/\lambda}$, $\hat{N}_t = N_t/(A_t/$

[6] The exponents and divisors in the second line of the learning transition equation in (2.4) keep the learning from becoming easier as the firm's size expands with rising demand and improving technology. As will be clear when the problem is expressed in stationary form, this means the firm faces the "same" learning problem with every technology.

[7] It is straightforward to show that (1) if labor input grows any faster, profits diverge to minus infinity, and (2) if labor input grows any slower, it falls below the level warranted by the static marginal product of labor.

$T_t)^{\epsilon-1}$, $\hat{w} = w_t/T_t$, and $\hat{Y}_{t-1} = Y_{t-1}/(A_t^{\epsilon}T_t^{1-\epsilon})$, the firm's problem is *equivalent* to[8]

$$\max_{\{\tau_t, \hat{N}_t\}} E_0 \sum_{t=0}^{\infty} \beta^t \left[\left(\hat{Q}_t^{\lambda} \hat{N}_t \right)^{1-1/\epsilon} D_t^{1/\epsilon} - \hat{w}\hat{N}_t \right]$$

subject to

$$D_t = D_{t-1}^{\rho} e^{\omega_t}, \tag{3.1}$$

$$\hat{Q}_t = \begin{cases} 1, & \text{if } \tau_t = 0, \\ \min\left(\gamma e^{-\mu\tau}, \hat{Q}_{t-1} + \hat{Y}_{t-1}\right), & \text{if } \tau_t = \tau_{t-1} + 1. \end{cases} \tag{3.2}$$

Therefore $\hat{Q}_t \in [1, \gamma)$. The firm can, at worst, adopt the leading-edge technology and operate it with productivity 1. The firm cannot operate at efficiency γ with the *current* leading-edge technology since it can only improve on 1 by accumulating experience.

The problem can be cast as the dynamic program

$$\max_{\{\tau_t, \hat{N}_t\}} E_t \left\{ \mathrm{v}\left(\tau_{t-1}, D_t, \hat{Q}_t\right) = \left(\hat{Q}_t^{\lambda}\hat{N}_t\right)^{1-1/\epsilon} D_t^{1/\epsilon} - \hat{w}\hat{N}_t + \beta\, \mathrm{v}\left(\tau_t, D_{t+1}, \hat{Q}_{t+1}\right) \right\} \tag{DP}$$

subject to (3.1) and (3.2). The state variables are (1) the vintage τ of last period's production technology, (2) the current level of demand D, and (3) the level of productivity the firm will enjoy *if* it continues to use last period's vintage of technology, $\hat{Q}$. The firm can adopt the leading-edge production technology by setting $\tau_t = 0$. Economic logic gives us an upper bound on the ergodic set for the age of the technology τ. A necessary condition for sticking with a technology is that $\gamma \exp(-\mu\tau) \geq 1$; if τ is such that this does not hold, then operating the leading-edge technology with no experience yields higher productivity than sticking with the old technology. The leading-edge technology is advancing at rate μ, so adopting it eventually dominates any given older vintage. Hence $\tau \in \{0, 1, \ldots, \text{Integer}(\ln(\gamma)/\mu)\}$.

A two-state approximation to (3.1) has $\ln(D) \in \{-\sigma, \sigma\}$ with $\sigma = \sigma_{\omega}/(1-\rho^2)^{1/2}$ and transition probabilities equal to $(1-\rho)/2$. Since $\hat{Q}_t \in [1, \gamma)$, I consider $\hat{Q}_{t+1} \in \{1, 1+0.001, 1+0.002, \ldots, \gamma - 0.001\}$, or productivity increments of roughly one-tenth of 1%.

Given the discrete state space, it is straightforward to iterate on (DP), starting from an initial guess v_0 for the value function. When the func-

[8] $\beta = \exp(\mu_e + (\mu - \mu_e)(\epsilon - 1))/(1+r)$. I omit a constant multiplying the present value of net cash flows.

TABLE I
Parameters

ϵ	Price elasticity of demand for the firm's output
ρ	First order autocorrelation of the stationary demand shock process
σ_ω	Standard deviation of innovations to the stationary demand shock process
μ_e	Deterministic growth rate of wages and demand
μ	Deterministic growth rate of the leading edge technology available to the firm
r	Real interest rate
γ	Ratio of maximal to initial productivity
λ	Elasticity of productivity with respect to experience
$\hat{w}$	Wage (scales the firm)

tional iterates converge within some tolerance, we have a value function and associated policy rules that approximately solve (DP). For the parameter values considered below, $\beta < 1$ so that discounted cash flow is finite. For every set of parameter values discussed below, I found that the value function converged within a tolerance of 0.001%.

4. CALIBRATION

Table I reviews the structural parameters of the model for easy reference. Table II gives the baseline set of parameter values, chosen as follows:

$\epsilon =$ *price elasticity of demand for the firm's output*. With production linear in labor, labor grows at the rate $(\epsilon - 1)(\mu - \mu_e)$, where ϵ is the price elasticity of demand. In the model μ is the exogenous growth rate of

TABLE II
Baseline Parameter Values

$\epsilon = 3.00$
$\rho = 0.95$
$\sigma_\omega = 0.018$
$\mu_e = 0.63\%$
$\mu = 0.63\%$
$r = 2\%$
$\lambda = 0.32$
$\gamma = 2.00$
$\hat{w} = 1.32$

industry total factor productivity (TFP). Thus, according to the model, a regression of average growth in N on the average growth rate of TFP across industries yields a consistent estimate of $(\epsilon - 1)$. Using 449 four-digit industries in the NBER Manufacturing Productivity Database [10], I regressed average employment growth on average TFP growth over the period 1960–1990. The OLS estimate of ϵ is 3.03 (standard error 0.39). Since $\hat{\epsilon}$ significantly exceeds 1, the estimate provides evidence that the demand for U.S. manufactures is price-elastic.

ρ, σ_ω, and μ_e. Recall that aggregate demand C can be constructed from

$$C_t = \int_0^M Y_t(i)^{1-1/\epsilon}\, di$$

and broken into stochastic stationary and nonstochastic trend components

$$C_t = D_t T_t,$$

$$D_t = D_{t-1} e^{\omega_t},$$

$$T_t = T_{t-1} e^{\mu_e}.$$

Using $\hat{\epsilon} = 3$ and annual data for the 449 manufacturing industries in the NBER Manufacturing Productivity Database, I constructed an annual version of C. Its properties were very similar to those of the simple sum of the 449 industries, or annual real manufacturing output. So for estimating ρ I consider the properties of quarterly real manufacturing output over the period 1960:–1990:4. I use OLS to estimate

$$\ln C_t = \text{constant} + \mu_e(1 - \rho)t + \rho \ln C_{t-1} + \omega_t.$$

The resulting estimates are $\mu_e = 0.0063$ (standard error 0.0031), $\rho = 0.95$ (standard error 0.03), and $\sigma_\omega = 0.018$.

μ. Although this is a partial equilibrium model, in general equilibrium the growth rate of consumption will equal the average growth rate of industry TFP. For this reason I set the baseline value for μ to be 0.63% per quarter, the same as μ_e. Of course, industry TFP growth rates do differ, so it is reasonable to consider individual industries as having $\mu \neq \mu_e$. Averaging TFP growth over the period 1960–1990 yields estimates of μ for each of the 449 manufacturing industries. Annual rates of TFP growth range from −3.6% per year for "Primary Smelting and Refining of Nonferrous Metals, Except Copper and Aluminum" (SIC 3339) to +9.0% per year for "Electronic Computing Equipment" (SIC 3573).

r. Typical estimates for the United States are a capital share of 1/3, a capital−output ratio of 2, and an annual depreciation rate of around 9%, suggesting an annual r of about 8%, or about 2% per quarter.

λ. The learning by doing literature is filled with estimates of learning clustered around the rate of 20%.[9] Many of these studies, such as Baloff [5−8], Garg and Milliman [21], and Irwin and Klenow [28], are for a single production process, just as required to calibrate the model here.[10] The 20% figure refers to the rate at which productivity rises with a doubling of cumulative output. Thus $2^{\lambda} = 1.2$, which implies $\lambda = 0.32$.

γ. Jovanovic and Nyarko [30] report "progress ratios"—ratios of peak to initial productivity—from a dozen empirical studies of learning by doing. The range is 1.14−2.9 for 10 of their 12 studies, so I consider $\gamma = 2$ as the baseline value. Since the range is rather wide, I also consider 1.5 and 2.5.

$\hat{w}$. The (detrended) wage determines the size of the firm and therefore the scale of the learning problem. In essence, units matter since experience starts at "1" and rises with cumulative output toward γ. I set $\hat{w}$ so that, conditional on the other parameter values, there is one technology update every 23 quarters, which is the average duration from peak to trough of postwar U.S. business cycles (see www.nber.org).

5. CHARACTERIZING TECHNOLOGY UPDATES

Table II lists the baseline set of parameter values. Table III characterizes optimal firm behavior under various sets of parameter values. Row 1 of Table III gives results with the baseline set, except with $\sigma_{\omega} = 0$ so that there are no demand shocks. Under these parameter values, the firm finds it optimal to update its technology every 23 quarters (as set in the scaling of the wage). The path of firm productivity over 100 quarters is shown in Fig. 1. Productivity falls 19% in update quarters. The firm could wait until the frontier technology is so much better than the old technology that productivity continues to rise through the shift to each new technology, but it is not optimal to do so. This is because rapid initial learning justifies earlier adoption of the new technology. If firms were staggered across different vintages, the ratio of peak-to-trough productivity across firms would be 1.24.

[9] See Yelle [35] and Argotte and Epple [1] for surveys.

[10] Bahk and Gort [3] estimate much slower learning by doing, but they do not control for changes in production technology at a plant.

TABLE III
Optimal Updating

	Frequency	TFP drop (%)	Peak TFP / Trough TFP
Baseline (with no shocks)	23	19	1.24
Faster frontier growth (μ = 3.5 vs. 2.5%)	19	18	1.23
Slower frontier growth (μ = 1.5 vs. 2.5%)	27	20	1.25
Bigger progress ratio (γ = 2.5 vs. 2)	32	26	1.32
Smaller progress ratio (γ = 1.5 vs. 2)	14	11	1.13
Faster learning (λ = 0.40 vs. 0.32)	28	22	1.28
Slower learning (λ = 0.20 vs. 0.32)	23	11	1.14
Higher real interest rate (r = 3 vs. 2%)	24	19	1.24
Lower real interest rate (r = 1 vs. 2%)	22	19	1.24
Higher elasticity of demand (ϵ = 3.5 vs. 3)	24	19	1.24
Lower elasticity of demand (ϵ = 2.5 vs. 3)	23	19	1.24

To gain insight into the model's properties, I consider alternative parameter values to see how they alter the model's implications. In the first exercise I increase μ, the speed at which each successive vintage improves on the previous vintage, from 2.5 to 3.5%. I find that updates are more frequent when frontier technology advances more rapidly. Specifically, when μ = 3.5% per year the firm updates every 19 quarters, as opposed to every 23 quarters when μ = 2.5% (see Fig. 2). For a given vintage of the technology in use, the experience lost when updating is no larger while the frontier technology adopted is better because of more rapid frontier advance. In other words, the opportunity cost of updating has not changed while the benefit has risen at all vintage levels. Thus more frequent updating. I also find that the productivity drop accompanying updates is

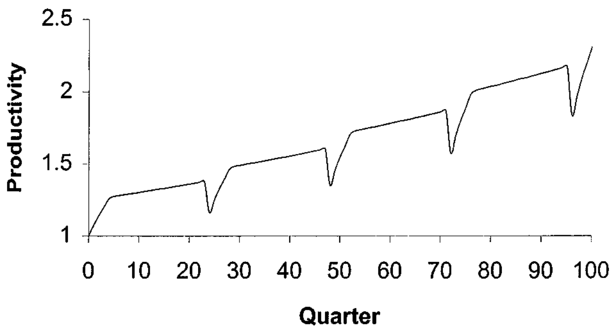

FIG. 1. Optimal updating.

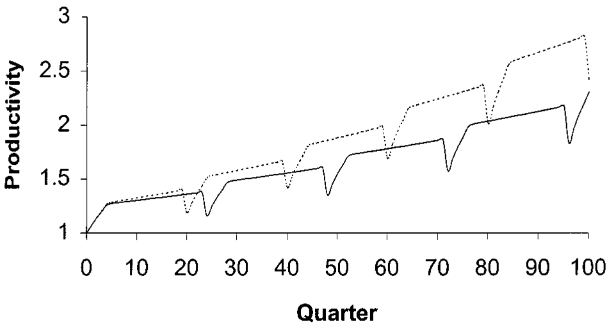

FIG. 2. Faster frontier growth.

slightly smaller (18 vs. 19%). The shorter interval over which each vintage is used means less experience is acquired with each vintage. The frequently updating firm will accept a smaller initial drop in productivity since it will not stick with the technology as long. The peak-to-trough ratio is slightly smaller with faster growth (1.23 vs. 1.24).[11]

In the second exercise I increase γ, the "progress ratio" or ratio of maximum to initial productivity with a technology (see Fig. 3). A higher γ means there is more to learn with each vintage. I consider $\gamma = 2.5$, compared to the baseline value of 2.0. The higher γ leads the firm to update less frequently. The firm sticks with a vintage for 32 quarters as opposed to 23. Since there is more to learn, the firm sticks with each technology longer and learns more with each. Moreover, both the productivity drop (26 vs. 19%) and the ratio of peak-to-trough productivity (1.32 vs. 1.24) rise. The firm tolerates a bigger drop in productivity because it will reap benefits of the new technology for a longer stretch. As Table III shows, the results are symmetric with respect to lowering γ to 1.5: higher frequency of updates, smaller productivity drops, and smaller peak-to-trough ratio.

[11] This prediction contrasts with Dwyer's [20] finding that the dispersion of plant productivity is larger in those U.S. textile industries with more rapid productivity growth. It is possible that the progress ratio γ tends to be higher in faster growing industries, which we will see would imply a higher peak-to-trough ratio. Where technology is changing rapidly, there may be more to learn about each technology. Consistent with this hypothesis, the ratio of peak-to-trough "yield" of chips is reputed to be about 9 (see Irwin and Klenow [28] for evidence that the chip defect rate goes from 90 down to 10% over the typical chip plant's life). This ratio is much higher than for other processes documented by Jovanovic and Nyarko [30], for which the technology is presumably changing more slowly than for semiconductors.

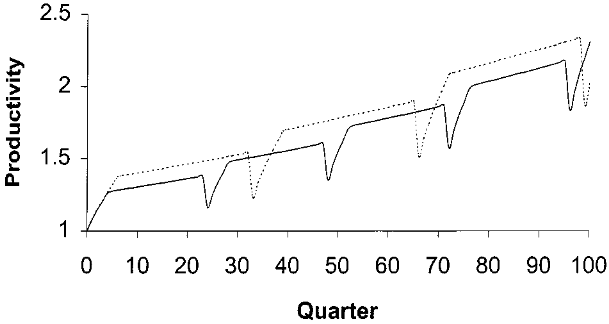

FIG. 3. Bigger progress ratio.

I next increase λ, the speed of learning by doing. I find that updates occur every 28 quarters when $\lambda = 0.40$, compared to every 23 quarters when $\lambda = 0.32$. The TFP drop and peak-to-trough TFP are higher at 22% (vs. 19%) and 1.28 (vs. 1.24). The higher learning rate encourages firms to stick with the old technology longer, and so they tolerate a larger drop. As Table III shows, the results are not symmetric with respect to lowering λ. When I let $\lambda = 0.2$, the firm uses each technology for 23 quarters, just like when $\lambda = 0.32$, but learns less and so exhibits a smaller TFP drop (11%) and a smaller peak-to-trough ratio (1.14). Note that in the extreme case of $\lambda = 0$, where there is no learning, the firm simply updates every quarter since there is no experience to lose.

The next two rows of Table III report results with higher and lower real interest rates (3 and 1%, respectively) than in the baseline case (2%). As expected from the investment nature of technology updating, the firm updates less frequently the higher is the interest rate. But the effect on timing is modest, which is not surprising given that the magnitude of the real interest rate is swamped by the magnitude of movements in productivity around updates.

The final two rows of Table III give the results with a higher and lower elasticity of product demand (3.5 and 2.5) than the baseline case (3). As shown, the results are not very sensitive to this parameter's value.

I now consider the impact of demand shocks on firm updating. For each of the two demand states (low demand and high demand), Table IV gives the mean age of the technology abandoned during that state and the fraction of technology updates occurring during that state. Under the estimated autocorrelation coefficient of $\rho = 0.95$ and the estimated shock standard deviation of $\sigma_\omega = 0.018$, firms update their technology after 25.0

TABLE IV
Optimal Updating with Demand Shocks

	Low demand	High demand
No shocks ($\rho = 0$, $\sigma_\omega = 0$)	23.0 (50%)	23.0 (50%)
Point Estimates ($\rho = 0.95$, $\sigma_\omega = 0.018$)	25.0 (46%)	23.1 (54%)
More persistence ($\rho = 0.99$, $\sigma_\omega = 0.018$)	26.0 (47%)	23.0 (53%)
Less persistence ($\rho = 0.50$, $\sigma_\omega = 0.018$)	23.2 (63%)	24.0 (37%)
No persistence ($\rho = 0$, $\sigma_\omega = 0.018$)	23.3 (75%)	24.0 (25%)
Bigger shocks ($\rho = 0.95$, $\sigma_\omega = 0.03$)	25.0 (46%)	23.1 (54%)
Smaller shocks ($\rho = 0.95$, $\sigma_\omega = 0.01$)	24.0 (48%)	23.0 (52%)

quarters (on average) if demand is low and after 23.1 quarters (on average) if demand is high. So firms delay technology updates in recessions, preferring to do them during booms (54% of all updates being carried out during high demand states). The next row illustrates that firms delay even more in recessions relative to booms when demand shocks are more persistent ($\rho = 0.99$ vs. 0.95).

In contrast, Cooper and Haltiwanger's [16] model of machine replacement predicts that firms will retool during slow times. They find that, consistent with their model, U.S. auto firms retool their plants in months of low seasonal demand. There are two differences between our models that together explain the divergence in our predictions. The first difference is that my model is one of protracted learning by doing, whereas theirs is one of a single month of downtime. In the presence of learning by doing, peak demand has the beneficent by-product of pushing the firm down the learning curve faster. When there is instead "learning by waiting," as in Cooper and Haltiwanger's setup, updating in recessions is preferred because the low transitional productivity is endured with a lower scale of production.

The second difference is that the current model focuses on business cycles rather than on seasonal cycles. Business cycles are more persistent than seasonal cycles, and the sensitivity of technology updates to demand shocks depends on their persistence. With prolonged learning a prolonged boom offers the benefit of faster learning. When shocks are not persistent in my model, I get Cooper and Haltiwanger's result that firms prefer to switch in slow times. As Table IV shows, when there is much less persistence ($\rho = 0.5$) or no persistence ($\rho = 0$) of the shocks, firms are more likely to update their technology in recessions (63 and 75% of updates, respectively, for the two values of ρ).

6. COMPARING THE MODEL'S IMPLICATIONS TO THE DATA

Industry Data

It is well known that aggregate labor input and productivity are procyclical (labor and productivity growth are each significantly positively correlated with real output growth). I find that this also holds true for most of the 449 manufacturing industries in the NBER Manufacturing Productivity Database over the period 1960–1990. The 1960–1990 correlations between growth in hours worked and output growth for each industry average 0.72 across the 449 industries. They have a standard deviation of 0.18 and range from −0.03 to 0.98. The correlations between TFP growth and output growth average 0.75, have a standard deviation of 0.14, and range from 0.14 to 0.96. A less well-known empirical regularity is a low correlation between hours growth and productivity growth. Christiano and Eichenbaum [13] find correlations of −0.20 and 0.16 between aggregate productivity and household and establishment survey hours. This pattern also holds in most of the 449 manufacturing industries. In 443 of the 449 industries the productivity–hours correlation is lower than the productivity–output correlation. The productivity–hours correlations average 0.40, while the productivity–output correlations average 0.75.

Christiano and Eichenbaum stress that technology-shock driven models imply productivity–hours correlations above 0.90 with standard errors around 0.10. These authors push the correlation toward zero by adding government consumption shocks, which affect hours without affecting TFP. They meet with modest success, pushing the correlation down to 0.57. One could drive the TFP–hours correlation still lower by adding more demand and input price shocks. Yet models with variable work effort (e.g., Bils and Cho [11]) and variable capital utilization (e.g., Bils and Klenow [12]) imply that demand and input price shocks move *measured* productivity in the same direction as hours. Adding an impulse or propagation mechanism that pushes hours and productivity in *opposite* directions seems necessary to fit the facts.

The sequence of bounded learning curves modeled here represents a strong mechanism pushing hours and productivity in opposite directions. When a firm updates its technology, its productivity falls and its labor input rises. Labor input rises because, although the "static" marginal product of labor is low, the "dynamic" marginal product is high because of the healthy gains in future productivity arising from the initial steepness of the learning curve.[12] As the firm accumulates experience with the vintage,

[12] By "static" is meant the current period marginal revenue product of labor. By "dynamic" is meant the static plus the discounted present value of resulting future gains in the marginal revenue product of labor.

productivity rises but labor input falls. Labor input falls because the learning curve flattens. In short, learning by doing separates the ‘‘static’’ marginal product of labor (which is highly correlated with TFP) from the ‘‘dynamic’’ marginal product of labor (which incorporates learning).

Table V provides correlations for the baseline model (with parameter values given in Table II). Panel A of Table V is for the case with no demand shocks, and panel B is for the case with the estimated demand shocks. The first entry shows that growth of labor input and growth of output are highly positively correlated (0.82 and 0.79 in the respective panels). In updating quarters, despite the TFP drop, labor input and output rise as firms take advantage of the steep learning curve. As experience accumulates, labor input and output fall, and TFP rises. Labor input is almost the mirror image of TFP (correlations −0.86 and −0.89 in the respective panels). The correlations between output growth and TFP growth are also negative (−0.44 in both panels), but substantially less so than the correlations between labor growth and TFP growth. The model matches the pattern of lower productivity−hours correlations than productivity−output correlations. The levels of the model correlations are too low, however, relative to their industry averages counterparts, which are 0.40 and 0.75. But the model here has no (within-technology) productivity shocks; adding them would naturally push the correlations upward. Adding variable utilization of labor would push the correlations up further. Combined with productivity shocks and variable utilization, a sequence of bounded learning curves could generate lower, more realistic correlations between hours and productivity.

Plant Data

A number of recent studies, such as Dhrymes [17], Baily, Hulten, and Campbell [4], Bartelsman and Dhrymes [9], and Dwyer [20], document features of U.S. manufacturing plants in the Census Longitudinal Re-

TABLE V
Correlations between Output, Labor, and TFP

	$\Delta \ln(Y)$	$\Delta \ln(N)$
A. Without demand shocks		
$\Delta \ln(N)$	0.82	
$\Delta \ln(\text{TFP})$	−0.44	−0.86
B. With demand shocks		
$\Delta \ln(N)$	0.79	
$\Delta \ln(\text{TFP})$	−0.44	−0.89

search Database. In interpreting their findings, it is important to be conscious of sample selection. Suppose, for example, that the samples are unbalanced (plants entering and exiting), that productivity is fixed at each plant, and that plants with higher productivity have lower exit rates. Then, as in Jovanovic [29], plant productivity rises with plant age even though there is no learning whatsoever. In the following I will therefore stress when the facts apply to continuously operating plants (i.e., balanced panels). Balanced panels entail selection of continuously surviving plants, however, and long-time survivors might have rising productivity relative to exiters and entrants. Clearly, the model in this paper does not adequately address exiters or entrants and how their behavior differs from continuously operating plants.[13]

Given the model's feature that only one technology can be used at a time, the plant seems to be the appropriate unit of empirical observation. Although one might be skeptical that existing plants adopt new technologies, many new technologies do not require building a new plant from scratch. In support of this claim, Dunne [19] finds that plant age and technology use (measured by the Census Survey of Manufacturing Technology) are virtually uncorrelated in U.S. manufacturing, and Doms and Dunne [18] document a pattern of ‘‘investment spikes’’ or bunching of equipment investment in 1 or 2 years of a 16-year span of a manufacturing plant's life, suggestive of incumbent plants updating their technologies.

I now focus on Baily et al.'s [4] study of Census year data (1963, 1967, 1972, 1977, and 1982) for 23 four-digit industries because its industry coverage is much broader than that of the other studies. For *stayers* (firms in the sample in adjacent Census years), some of the relevant Bailey et al. findings are:

1. ‘‘Plant-level fixed effects in productivity are persistent’’ over 5- and 10-year intervals. This is clearly *inconsistent* with the model, which has firms cycling through the productivity distribution. Thus the model would need to be modified to include permanent sources of plant heterogeneity alongside the updating.

2. ‘‘There is rather weak evidence for plant vintage effects.’’ This is consistent with the model, in which plants are updating their technologies so that old firms do not necessarily have older technology. As Baily et al. state, ‘‘Old plants invest in new equipment and acquire new technology

[13] The model's implication that firms update technology in booms (to learn rapidly) could help explain why entry rates are higher in booms (presuming entrants are using technology new to them). Learning by doing combined with uncertain demand might then explain entry waves, as documented, for example, by Gort and Klepper [25].

that way.'' Note that because the sample is of stayers, this result is not driven by the departure of young, unproductive plants.

3. ''Large offsetting effects on [industry] productivity come from the plants that are moving up rapidly in the distribution and from the plants that are falling rapidly'' over 5-year intervals. This is consistent with the model's prediction that, in the absence of full synchronization, firms will cycle in their productivity position—falling when they adopt a technology, rising as they gain experience operating it, then falling again as their technology becomes old relative to other firms.

Baily et al. also find that new plants have below-average productivity but tend to catch up. They report that part of this is because lower productivity plants exit, leaving the higher productivity entrants as the survivors, as in Jovanovic [29]. But they also report that surviving entrants' TFP-grows more rapidly than the average for incumbents. Presuming that new plants are using technology which is new to them, this finding is consistent with a startup period of low productivity followed by rapid learning.

Investment Spikes and the Procyclicality of Investment

In the model technology updates are not necessarily embodied in physical capital. But many changes in technology do involve new physical capital, such as those in the 1988 Survey of Manufacturing Technology conducted by the U.S. Census Bureau. This survey, described in Dunne [19], focused on new factory automation technologies such as programmable controllers, computer-automated design, and numerically controlled machines. In the case of information technology, equipment acquisition appears to be a necessary condition for technology adoption (see, for example, Yorukoglu [36]).

Although the model set out in Section 2 did not explicitly involve new capital, one could reinterpret some fraction of the TFP drop after technology updates as new investment, with TFP appropriately revised upward. For the interpretations to be truly isomorphic, the firm must invest more the earlier it adopts a technology of a given vintage. This certainly seems plausible in the case of computers, and for equipment more generally in light of the downward trend in the relative price of equipment (see Greenwood, Hercowitz, and Krusell [26]). Moreover, Bahk and Gort [3] estimate that, because of embodied technological change, each 1-year change in the vintage of equipment is associated with a 2.5–3.5% change in output.[14] This is in line with $\mu = 2.5\%$ here for the rate of frontier technology advance.

[14] See also Gort, Bahk, and Wall [24].

Under an equipment interpretation, the prediction of the model that updating occurs more often in booms than in recessions seems a decided advantage over models with updating concentrated in recessions. As is well known, investment is procyclical at the aggregate and industry level, and Doms and Dunne [18] document that investment spikes at the manufacturing plant level are highly procyclical. As discussed earlier, the key force in the model tilting updating toward booms is that learning is by *doing*. Learning by doing makes updating in periods of high demand attractive, as high output rates have the by-product of faster learning.[15] If, as in Cooper and Haltiwanger [16], learning required time rather than doing, firms would prefer to work out the bugs of new technologies in recessions rather than in booms.[16] Thus, along with considerable micro evidence, the strong procyclicality of investment points to learning by doing.

A related topic is Goolsbee's [22] study of aircraft retirement. He finds that airlines are more likely to retire aircraft in a recession. The assumption that retirement of old aircraft is synonymous with upgrading to new aircraft seems highly questionable, however. As Goolsbee and Gross [23] show, purchases of new aircraft are highly procyclical (and lumpy, as predicted by the model of technology updating here). In a typical recession firms increase retirement and *reduce* new aircraft purchases. Thus Goolsbee's evidence on aircraft retirements does not contradict the model here.

7. CONCLUSION

Micro evidence suggests that productivity gains from learning by doing diminish as experience accumulates with a technology. It also suggests that the learning may be largely specific to each production technology. I build these features into a model of a firm deciding when to update its technology. The model implies that the firm's productivity falls when it adopts a new technology, but grows quickly as the firm acquires experience with the new technology. The model further implies that firms accelerate technology adoptions during peak periods and delay them during trough periods. The intuition for this property of the model is that peak periods mean lots of doing and therefore lots of learning.

[15] Power [33] finds that plant productivity growth is rapid after investment spikes, consistent with a burst of learning after technology updating. She also finds a falling probability of an investment spike with years since the last spike. Controlling for μ, my model predicts a rising hazard rate. It is possible, however, that there are large differences in μ which generate a falling hazard rate. Moreover, in the data there are "spikelets" in years adjacent to spikes, which could result from a single spike straddling two calendar years.

[16] Parente [31] also features learning with time, rather than with cumulative output, so his model too would imply updating during recessions.

I find that the model may help explain features of U.S. manufacturing industries from 1960 to 1990, as well as some properties of plant level data reported in other studies. The model has a mechanism driving down the correlation between hours and productivity toward that observed in the data. It also implies that entrants have below-average productivity but quickly gain on the productivity of incumbent firms, again consistent with the evidence (with the evidence holding up even when confined to those plants which survive a long time). Finally, if the model is interpreted as involving substantial investment at the time of updating, it may help explain the procyclicality of investment at the aggregate, industry, and plant levels, including the investment spikes observed at manufacturing plants.

REFERENCES

1. L. Argotte and D. Epple, Learning curves in manufacturing, *Science* **247** (1990), 920–924.
2. H. Asher, "Cost–Quantity Relationships in the Airframe Industry," Report R-291, RAND Corporation, 1961.
3. B.-H. Bahk and M. Gort, Decomposing learning by doing, *J. Polit. Econ.* **101** (1993), 561–583.
4. M. N. Baily, C. Hulten, and D. Campbell, Productivity dynamics in manufacturing plants, *in* "Brookings Papers on Economic Activity: Microeconomics," pp. 307–358, Brookings Institute, Washington, D.C., 1992.
5. N. Baloff, Startups in machine-intensive production systems, *J. Industrial Eng.* **14** (1966), 25–32.
6. N. Baloff, Estimating the parameters of the startup model—An empirical approach, *J. Industrial Eng.* **18** (1967), 248–253.
7. N. Baloff, Startup Management, *IEEE Trans. Eng. Manag.* **EM-17** (1970), 132–141.
8. N. Baloff, Extensions of the learning curve—some empirical results, *Oper. Res. Quart.* **22** (1971), 329–340.
9. E. J. Bartelsman and P. J. Dhrymes, "Productivity Dynamics: U.S. Manufacturing Plants, 1972–1986," Finance and Economics Discussion Series 93-43, Board of Governors of the Federal Reserve System, 1994.
10. E. J. Bartelsman and W. Gray, "The NBER Manufacturing Productivity Database," Technical working paper 205, National Bureau of Economic Research, 1996.
11. M. Bils and J. O. Cho, Cyclical factor utilization, *J. Monetary Econom.* **33** (1994), 319–354.
12. M. Bils and P. J. Klenow, Using consumer theory to test competing business cycle models, *J. Polit. Econ.* **106** (1998), 233–261.
13. L. J. Christiano and M. Eichenbaum, Current real-business-cycle theories and aggregate labor-market fluctuations, *Amer. Econom. Rev.* **82** (1992), 430–450.
14. E. B. Cochran, New concepts of the learning curve, *J. Industrial Eng.* **11** (1960), 317–327.
15. R. Conway and A. Schultz, The manufacturing progress function, *J. Industrial Eng.* **10** (1959), 39–53.
16. R. Cooper and J. Haltiwanger, The aggregate implications of machine replacement: Theory and evidence, *Amer. Econom. Rev.* **83** (1993), 360–382.

17. P. J. Dhrymes, ''The Structure of Production Technology: Productivity and Aggregation Effects,'' Discussion Paper 91-5, Center for Economic Studies, U.S. Department of Commerce, Bureau of the Census, 1991.
18. M. Doms and T. Dunne, Capital adjustment patterns in manufacturing plants, *Rev. Econom. Dynam.* **1** (1998), 409–429.
19. T. Dunne, Plant age and technology use in U.S. manufacturing industries, *Rand J. Econom.* **25** (1994), 488–499.
20. D. W. Dwyer, Technology locks, creative destruction, and nonconvergence in productivity levels, *Rev. Econom. Dynam.* **1** (1998), 430–473.
21. A. Garg and P. Milliman, The aircraft progress curve—Modified for design changes, *J. Industrial Eng.* **12** (1961), 23–27.
22. A. Goolsbee, The business cycle, financial performance, and the retirement of capital goods, *Rev. Econom. Dynam.* **1** (1998), 474–496.
23. A. Goolsbee and D. Gross, ''Estimating the Costs of Adjustment: Heterogeneous Capital, Irreversibility, and Non-Convexity,'' mimeo, University of Chicago, 1997.
24. M. Gort, B.-H. Bahk, and R. A. Wall, Decomposing technical change, *Southern Econom. J.* **60** (1993), 220–234.
25. M. Gort and S. Klepper, Time paths in the diffusion of product innovations, *Econom. J.* **92** (1982), 630–653.
26. J. Greenwood, Z. Hercowitz, and P. Krusell, Long-run implications of investment-specific technological change, *Amer. Econom. Rev.* **87** (1997), 342–362.
27. W. Z. Hirsch, Manufacturing progress functions, *Rev. Econom. Statist.* **34** (1952), 143–155.
28. D. A. Irwin and P. J. Klenow, Learning-by-doing spillovers in the semiconductor industry, *J. Polit. Econ.* **102** (1994), 1200–1227.
29. B. Jovanovic, Selection and the evolution of industry, *Econometrica* **50** (1982), 649–670.
30. B. Jovanovic and Y. Nyarko, A Bayesian learning model fitted to a variety of empirical learning, curves, *in* ''Brookings Papers on Economic Activity: Microeconomics,'' pp. 247–305, Brookings Institute, Washington, D.C., 1995.
31. S. Parente, Technology adoption, learning by doing and economic growth, *J. Econom. Theory* **63** (1994), 346–369.
32. C. C. Pegels, On startup or learning curves: An expanded view, *AIIE Trans. Eng. Manag.* **1** (1969), 216–223.
33. L. Power, ''The Causes and Consequences of Investment Spikes in U.S. Manufacturing Plants,'' Ph.D. dissertation, University of Maryland, 1994.
34. J. H. Russell, Progress function models and their deviations, *J. Industrial Eng.* **19** (1968), 5–11.
35. L. E. Yelle, The learning curve: Historical review and comprehensive survey, *Decision Sci.* **10** (1979), 302–328.
36. M. Yorukoglu, The information technology productivity paradox, *Rev. Econom. Dynam.* **1** (1998), 551–592.
37. A. Young, Invention and bounded learning by doing, *J. Polit. Econ.* **101** (1993), 443–472.
38. R. Zeckhauser, Optimality in a world of progress and learning, *Rev. Econom. Stud.* **35** (1968), 363–365.