Federal Reserve Bank of Minneapolis
Research Department Staff Report 296

Revised September 2003

The Transition to a New Economy After the Second Industrial Revolution*

Andrew Atkeson

University of California, Los Angeles,
Federal Reserve Bank of Minneapolis,
and National Bureau of Economic Research

Patrick J. Kehoe

Federal Reserve Bank of Minneapolis,
University of Minnesota,
and National Bureau of Economic Research

ABSTRACT

During the Second Industrial Revolution, 1860–1900, many new technologies, including electricity, were invented. After this revolution, however, several decades passed before these new technologies diffused and measured productivity growth increased. We build a quantitative model, in which technology diffusion requires learning, that we use to study this transition to a new economy. Our main innovation is our method of using micro data on the patterns of birth, growth, and death of plants in the U.S. economy to quantify the parameters of the learning process. Our model generates both slow diffusion and a delay in growth after the Second Industrial Revolution similar to that in the data. Our model casts doubt, however, on the conjecture that this experience is a useful parallel for understanding the productivity paradox after the Information Technology Revolution.

*Atkeson and Kehoe thank the National Science Foundation. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

The period 1860–1900 is often called the *Second Industrial Revolution* because of the large number of technologies invented during that time. Historians such as Schurr et al. (1960), Rosenberg (1976), Devine (1983), and David (1990, 1991) focus on the development of new technologies based on electricity that began in the Second Industrial Revolution as the driving force behind the rapid technical change in manufacturing that continued throughout the 20th century. This revolution led eventually to a *new economy,* characterized by faster growth in productivity, as measured by output per hour. Following this revolution, however, there was a delay of several decades before these new technologies diffused widely among manufacturing plants and this new economy appeared. David (1990) refers to this slow transition to a new economy after the Second Industrial Revolution as a *productivity paradox.* This slow transition is indeed paradoxical from the viewpoint of the standard growth model. In this model, technologies are disembodied, new technologies diffuse instantly, and faster technical change results immediately in faster growth of measured productivity.

We evaluate the ability of hypotheses put forward by David (1990) and other historians to quantitatively explain the slow transition to a new economy that followed the Second Industrial Revolution. We do so by building a quantitative model of technology diffusion intended to capture the main elements of the historians' hypotheses, including David and Wright's (1999, p. 4) hypothesis that "the slow pace of adoption [of electric power] prior to the 1920s was largely attributable to the unprofitability of replacing still serviceable manufacturing plants embodying production technologies adapted to the old regime of mechanical power derived from water and steam."

Our model builds in three key assumptions motivated by the hypotheses of these historians.

- *New plants embody new technologies.*[1] This assumption is motivated by the work of Devine (1983) and David (1990, 1991), who argue that manufacturing plants needed to be completely redesigned in order to make good use of the new technologies stemming from the development of electric power.
- *Improvements in the technology for new plants are ongoing.* Specifically, we model the transition to a new economy after the Second Industrial Revolution as arising from a once-and-for-all increase in the rate of improvement in the frontier technology embodied in the design of new plants. This assumption is meant to capture the arguments of Devine (1990) and Sonenblum (1990) that the process of improving efficiency through changes in factory design after the Second Industrial Revolution continued for decades, through at least the 1980s, and lay behind the new economy after this revolution.
- *New plants improve their technology through an extended period of learning.* This assumption is consistent with a broad body of work on learning as well as the discussion in Chandler (1992). We use micro data on plants to argue that substantial learning at the plant level continues for at least 20 years.

(For a more detailed discussion of the links between the historical analyses and our model, see Appendix A.)

We pose this quantitative question: When the parameters of our model are set to match micro data on the learning process at the plant level, can the model generate the observed slow diffusion of electricity and the slow transition to a new economy after the Second Industrial Revolution? We find that it can. The model generates the slow improvement in productivity

growth because agents slowly abandon plants embodying old technologies in which they have built up knowledge for new plants embodying new technologies that will be fully productive only after an extended period of learning.

Our main conceptual innovation here is our method of using micro data on the life cycle patterns of plants in the U.S. economy—their birth, growth, and death—to uncover the parameters of the learning process at the plant level. Our method is motivated by two observations from the micro data. One is that the size of plants—as measured by the share of the manufacturing labor force employed in each new cohort of manufacturing plants—tends to grow substantially for at least 20 years. The other is that the average productivity of labor and capital in these plants relative to that of other plants neither rises nor falls substantially with age.

We propose the simplest model of learning that matches these two observations. We model learning as changes in plant-specific productivity. We allow labor and capital to flow to equate marginal products across plants, and we assume Cobb-Douglas production. In our model, as a plant's specific productivity increases, the plant grows by adding labor and capital so as to keep the marginal products of its factors equal to those of other plants. Hence, the observation that the employment share in a cohort of plants grows substantially for at least 20 years implies that the average of their specific productivities does as well. We thus infer that the amount of learning is substantial and that learning extends at least over this time period. The Cobb-Douglas assumption implies that, even though plants undergo substantial growth in their specific productivities as they age, they show no growth in their average labor and capital productivity relative to other plants. This leads our model to be consistent with the observation that the average productivity of plants neither rises nor falls substantially

with age. (The idea of linking specific productivity and size is familiar from Hopenhayn and Rogerson 1993.)

Our method of measuring learning at the plant level contrasts sharply with the standard approach in the learning literature, exemplified by the work of Bahk and Gort (1993). These researchers associate learning in a plant with changes in the average productivity of labor and capital at that plant relative to other plants. In the context of our model, their approach is conceptually flawed. Indeed, if it were applied to our model, it would find no learning at all, regardless of how much learning was actually going on. Our approach to measuring learning is one of the features that distinguish our work from related models of transition based on embodiment and learning, including those of Hornstein and Krusell (1996) and Greenwood and Yorukoglu (1997).

When the parameters of our model are set using our method, the model's predictions match the U.S. data surprisingly well. The model's path of diffusion of new embodied technologies is similar to the data's path of diffusion of electric power in U.S. manufacturing plants during 1869–1939. The model's trends in the growth of output per hour are similar to the data's for 1869–1969. We find it intriguing that the model matches these data so well even though we did not attempt to replicate any features of the transition after the Second Industrial Revolution when we chose the parameters of the learning process.

We then demonstrate that all of our three key assumptions are necessary for our model to generate the slow diffusion and slow transition to a new economy. We show that if most of the increase in the pace of technical change comes from new technologies that are not embodied in plants but rather are disembodied, as they are in the standard growth model, then the transition is almost immediate. To generate a new economy with a higher

growth rate of output, we must assume that the increase in the pace of technical change is ongoing. And without learning at the plant level, the model generates a rapid diffusion of new technologies and a rapid transition.

Having shown that the model can account for the slow transition after the Second Industrial Revolution, we use the model to study David's (1990) additional conjecture that this experience is a useful parallel for understanding the slow transition after the more recent Information Technology Revolution. We find that it is not. The model predicts, counterfactually, that the transition after the Information Technology Revolution should be fast. This is because in the model, when the pace of technical change is relatively fast in an old economy, as the data show was true before the recent revolution, the diffusion of new technologies is also fast, and manufacturers build up only a relatively small amount of knowledge about a given technology. At the start of a transition from such an old economy, therefore, manufacturers are willing to quickly abandon this small stock of knowledge in order to adopt a superior technology. Of course, in the data, the transition appears to have been quite slow. Apparently, another approach is required in order to explain the experience after the Information Technology Revolution. In our conclusion, we discuss how one might modify the model to account for this more recent productivity paradox.

Our study is related to several strands of literature that we discuss after we present the model and the results.

1. Diffusion and Productivity Growth After the Second Industrial Revolution

Many of the technologies that had a profound impact on living standards in the 20th century were invented between 1860 and 1900. These technologies include electricity, the internal combustion engine, petroleum and other chemicals, telephones and radios, and indoor plumbing. (For a description of technological inventions during this time, see Gordon 2000a.) While all of these inventions undoubtedly had a substantial economic impact, we follow Schurr et al. (1960, 1990), Rosenberg (1976), Devine (1983), and David (1990, 1991) and focus on the new technologies based on electricity.

In this section, we document the gradual diffusion of electric power in U.S. manufacturing over the period 1869–1939 and the gradual increase in the growth of productivity—output per hour—in U.S. manufacturing over the period 1869–1969. (We choose 1869 as the starting point because the early data are derived from the U.S. Census Bureau's censuses of manufacturing establishments, which are taken every decade starting in 1869.)

To document the slow diffusion of electricity, in Figure 1 we plot the fractions of mechanical power in U.S. manufacturing establishments that are derived from water, steam, and electricity during 1869–1939. (These data are from Table 3 in Devine 1983.) Before 1899, more than 95% of mechanical power was derived from water and steam. Between 1899 and 1929, electricity use gradually replaced water and steam, so that by 1929, over 75% of mechanical power was electric. If we measure the diffusion of electricity starting in 1869, then we see that it took 50 years for electricity to provide 50% of mechanical power. This measure of the speed of diffusion is sensitive to the choice of starting date. A measure of the speed of diffusion that is less sensitive to that choice is the time required for a technology to diffuse

from 5% to 50%. For electricity in U.S. manufacturing, this occurred in about 20 years, from 1899 to 1919.

To document the consequent gradual increase in productivity growth, in Figure 2 we plot annual output per hour in U.S. manufacturing over the period 1869–1969. (These data are from the U.S. Department of Commerce 1973.) We also show in this figure linear trends for the three periods 1869–99, 1899–1929, and 1949–69. (These periods are chosen to omit the Great Depression and World War II.) The trend growth rates of output per hour in these three periods increase gradually, from 1.6% to 2.6% to 3.3%. (Gordon 2000b documents a similar gradual acceleration for the growth of output per hour for the U.S. economy as a whole.)

2. A Model of Technology Diffusion

Here we describe the model we use to generate a transition to a new economy. In the model, time is discrete and is denoted by periods $t = 0, 1, 2, \ldots$. The economy has two types of agents: workers and managers. There exist a continuum of size 1 of workers and a continuum of size 1 of managers.

Workers are each endowed with one unit of labor per period, which they supply inelastically. Workers are also endowed with the initial stock of physical capital and ownership of the plants that exist in period 0. Workers have preferences over consumption given by $\sum_{t=0}^{\infty} \beta^t \log(c_{wt})$, where β is the discount factor. Given sequences of wages and intertemporal prices $\{w_t, p_t\}_{t=0}^{\infty}$, initial capital holdings k_0, and an initial value a_0 of the plants that exist in period 0, workers choose sequences of consumption $\{c_{wt}\}_{t=0}^{\infty}$ to maximize utility subject to

the budget constraint

$$(1) \qquad \sum_{t=0}^{\infty} p_t c_{wt} \leq \sum_{t=0}^{\infty} p_t w_t + k_0 + a_0.$$

Managers are endowed with one unit of managerial time in each period. Managers have preferences over consumption given by $\sum_{t=0}^{\infty} \beta^t \log(c_{mt})$. Given sequences of managerial wages and intertemporal prices $\{w_{mt}, p_t\}_{t=0}^{\infty}$, managers choose consumption $\{c_{mt}\}_{t=0}^{\infty}$ to maximize utility subject to the budget constraint $\sum_{t=0}^{\infty} p_t c_{mt} \leq \sum_{t=0}^{\infty} p_t w_{mt}$. Notice that we have given all the initial assets to the workers. Since worker and manager utilities are identical and homothetic, aggregate variables do not depend on the initial allocation of assets.

Production in this economy is carried out in plants. In any period, a plant is characterized by its *specific productivity* A and its *age* s. We refer to the pair (A, s) as the plant's *organization-specific capital,* or simply its *organization capital.* This pair summarizes the built-up knowledge that distinguishes one organization from another.

To operate, a plant uses one unit of a manager's time, physical capital, and (workers') labor as variable inputs. If a plant with specific productivity A operates with one manager, capital k, and labor l, then the plant produces output

$$(2) \qquad y = zA^{1-\nu} F(k, l)^{\nu},$$

where the function F is linearly homogeneous of degree 1 and the parameter $\nu \in (0, 1)$. Notice that this technology implies that there are decreasing returns at each plant but constant returns for the economy as a whole. The technology parameter z is common to all plants and grows at an exogenous rate. We call z *economy-wide productivity.* Following Lucas (1978, p. 511), we call ν the *span of control* parameter of the plant's manager. The parameter ν may

be interpreted more broadly as determining the degree of diminishing returns at the plant level.

The timing of events in period t is as follows. The decision whether to operate or not is made at the beginning of the period. Plants that do not operate produce nothing; the organization capital in these plants is lost permanently. Plants with organization capital (A, s) that do operate, in contrast, hire a manager, capital k_t, and labor l_t and produce output according to (2). At the end of the period, operating plants draw independent multiplicative innovations, or *shocks,* ϵ to their specific productivity, with probabilities given by age-dependent distributions $\{\pi_s\}$. Thus, a plant with organization capital (A, s) that operates in period t has stochastic organization capital $(A\epsilon, s+1)$ at the beginning of period $t+1$.

Consider the process by which a new plant enters the economy. Before a new plant can enter in period t, a manager must spend period $t-1$ preparing and adopting a plan, or *blueprint,* for constructing the plant that determines the plant's initial specific productivity τ_t. Blueprints adopted in period $t-1$ embody the frontier technology regarding the design of plants at that point in time. These *frontier blueprints* evolve exogenously, according to the sequence $\{\tau_t\}_{t=0}^{\infty}$. Thus, a plant built in $t-1$ starts period t with initial specific productivity τ_t and organization capital $(A, s) = (\tau_t, 0)$.

We assume that capital and labor are freely mobile across plants in each period. Thus, for any plant that operates in period t, the decision of how much capital and labor to hire is static. Given a rental rate for capital r_t, a wage rate for labor w_t, and a managerial wage w_{mt}, the operating plant chooses employment of capital and labor to maximize static returns:

$$\max_{k,l} z_t A^{1-\nu} F(k,l)^{\nu} - r_t k - w_t l - w_{mt}. \tag{3}$$

The static returns to the owner of a plant with organization capital (A, s) in t are given by $d_t(A) - w_{mt}$, where $d_t(A) = z_t A^{1-\nu} F\big(k_t(A), l_t(A)\big)^{\nu} - r_t k_t(A) - w_t l_t(A)$ denotes static returns plus payments to managers and $k_t(A)$ and $l_t(A)$ are the solutions to this problem. Let $y_t(A) = z_t A^{1-\nu} F\big(k_t(A), l_t(A)\big)^{\nu}$ denote the associated output of the plant.

The decision whether or not to operate a plant is dynamic. This decision problem is described by the Bellman equation

$$(4) \qquad V_t(A, s) = \max\left[0, V_t^c(A, s)\right]$$

$$V_t^c(A, s) = d_t(A) - w_{mt} + \frac{p_{t+1}}{p_t} \int_{\epsilon} V_{t+1}(A\epsilon, s+1)\pi_{s+1}(d\epsilon),$$

where the sequences $\{\tau_t, w_t, r_t, w_{mt}, p_t\}_{t=0}^{\infty}$ are given. The value $V_t(A, s)$ is the expected discounted stream of returns to the owner of a plant with organization capital (A, s). This value is the maximum of the returns from closing the plant and those from operating it. The term $V_t^c(A, s)$, the expected discounted value of operating a plant of type (A, s), consists of current returns $d_t(A) - w_{mt}$ and the discounted value of expected future returns $V_{t+1}(A, s)$. The plant operates only if the expected returns $V_t^c(A, s)$ from operating it are nonnegative.

The decision whether or not to hire a manager to prepare a blueprint for a new plant is also dynamic. In period t, this decision is determined by the equation

$$(5) \qquad V_t^0 = -w_{mt} + \frac{p_{t+1}}{p_t} V_{t+1}(\tau_{t+1}, 0).$$

The value V_t^0 is the expected stream of returns to the owner of a new plant, net of the cost w_{mt} of paying a manager to prepare the blueprint for the plant. Managers are hired to prepare blueprints for new plants only if $V_t^0 \geq 0$. Since there is free entry into the business of starting new plants, in equilibrium we require that $V_t^0 \leq 0$. We summarize this condition

as $V_t^0 \phi_t = 0$, where $\phi_t \geq 0$ denotes the measure of managers preparing blueprints for new plants in t.

Let μ_t denote the measure at the beginning of period t of organization capital across plants that might operate in that period, where $\mu_t(A, s)$ is the measure of plants of age s with productivity less than or equal to A. Denote the measure of plants that operate in t by $\lambda_t(A, s)$. This measure is determined by μ_t and the sign of the function $V_t^c(A, s)$ according to

$$\lambda_t(A, s) = \int_0^A 1_{V^c}(a, s)\mu_t(da, s),$$

where $1_{V^c}(a, s) = 1$ if $V_t^c(a, s) \geq 0$ and 0 otherwise. For each plant that operates, an innovation to its specific productivity is drawn, and the distribution μ_{t+1} is determined from $\lambda_t, \phi_t, \{\pi_s\}$, and $\{\tau_t\}$ as follows:

$$\text{(6)} \qquad \mu_{t+1}(A', s+1) = \int_A \pi_{s+1}(A'/A)\lambda_t(dA, s)$$

for $s \geq 0$, where A' denotes specific productivity at $t+1$ and $\mu_{t+1}(\tau_{t+1}, 0) = \phi_t$.

Now let k_t denote the aggregate physical capital stock. Then the resource constraints for physical capital and labor are $\sum_s \int_A k_t(A)\lambda_t(dA, s) = k_t$ and $\sum_s \int_A l_t(A)\lambda_t(dA, s) = 1$. The resource constraint for aggregate output is $c_{wt} + c_{mt} + k_{t+1} = y_t + (1-\delta)k_t$, where δ is the depreciation rate of capital and y_t is defined by $y_t = \sum_s \int_A y_t(A)\lambda_t(dA, s)$. The resource constraint for managers is

$$\text{(7)} \qquad \phi_t + \sum_s \int_A \lambda_t(dA, s) = 1.$$

In equilibrium, the value of the workers' initial assets, or the initial value of plants, is $a_0 = \sum_s \int_A V_0(A, s)\mu_0(dA, s)$.

Given a sequence of frontier blueprints and economy-wide productivities $\{\tau_t, z_t\}$, initial endowments k_0 and a_0, and an initial distribution of organization capital across plants μ_0, an *equilibrium* in this economy is a collection of sequences of consumption and aggregate capital $\{c_{mt}, c_{wt}, k_t\}$; allocations of capital and labor across plants $\{k_t(A), l_t(A)\}$; measures of operating plants, potentially operating plants, and managers preparing blueprints for plants $\{\lambda_t, \mu_t, \phi_t\}$; value functions and operating decisions $\{d_t, V_t, V_t^c, V_t^0\}$; and prices $\{w_t, r_t, w_{mt}, p_t, \}$, all of which satisfy the above conditions.

In this equilibrium, the dynamic decision of whether to operate a plant has a simple form.

PROPOSITION 1. *In each period t, the decision to operate a plant is summarized by an age-dependent cutoff rule $A_t^*(s)$. In period t, plants of age s with specific productivity $A \geq A_t^*(s)$ continue operating, and those with $A < A_t^*(s)$ close.*

Proof. The proof follows from the result that the expected returns from operating a plant in period t, $V_t^c(A, s)$, are increasing in the plant's current productivity A for each age s. Clearly, current returns $d_t(A)$ are increasing in A. Also, since $A_{t+1} = A_t \varepsilon_t$, the distribution of a plant's productivity at $t+1$ is increasing in its productivity at t, in the sense of first-order stochastic dominance. The result that the value functions $V_{t+1}(A, s)$ are nondecreasing in A then follows. Thus, $V_t^c(A, s)$ is increasing in A for each age s. Letting $A_t^*(s)$ be defined by $V_t^c(A, s) = 0$, the result follows. Q.E.D.

To get a sense of the process for the birth, growth, and death—or the life cycle—of plants which our model generates, consider Figure 3. Here we show the evolution of the specific productivity of two plants that both enter in period t =1860. Both of these plants

start with productivity equal to that of the frontier blueprint in 1860, namely, τ_{1860}. This frontier blueprint grows exogenously over time at a constant rate as shown by the straight line labeled log τ_t. The two plants each experience random shocks to their plant-specific productivity drawn from age-dependent distributions π_s. Plant 1 is relatively lucky in that it draws especially favorable shocks to its specific productivity, while plant 2 is relatively unlucky.

In every period, each plant makes a decision whether to continue or to close, or *exit*. This decision is based on a comparison of the plant's current specific productivity and its future prospects for learning determined by π_s relative to the alternative of exiting and starting a new plant with the current frontier blueprint. The decision is summarized by an age-dependent cutoff rule as described in Proposition 1. Plant 1 has relatively high specific productivity; hence, it exits only after operating 30 years. In contrast, plant 2 has relatively low specific productivity; it exits much sooner. After each of these plants exits, the manager of the plant starts a new plant with the current frontier blueprint and begins the process of building up specific productivity in the new plant.[2]

In our model, technologies embodied in plants diffuse as new plants embodying these technologies are born and grow. Figure 3 also illustrates the mechanics of this diffusion. In 1863, the manager of plant 2 decides to exit and start a new plant that embodies the frontier blueprint of 1864 and then begins to learn with that new technology. Likewise, in 1890 the manager of plant 1 decides to exit and start a new plant that embodies the frontier blueprint of 1891 and then begins to learn with that new technology. In this manner, new plants embodying new technologies gradually replace old ones. Since our model has many such plants, each with different shocks to specific productivity, this diffusion of new embodied

technologies occurs smoothly over time.

Formally, we measure the diffusion of embodied technologies as follows. Let

$$y_{t,s} = \int_A \frac{y_t(A)}{y_t} \lambda_t(dA, s)$$

denote the fraction of total output y_t produced in plants of age s and

$$D_{t,t+k} = \sum_{s=0}^{k} y_{t,s} \tag{8}$$

be the fraction of total output y_t produced in plants of age k and younger. We measure the *diffusion in period $t+k$ of embodied technologies developed in period t or later* by the $D_{t,t+k}$, which, in the model, is the fraction of output produced in plants using technologies developed in period t or later.

3. Measuring Learning

The process governing learning is a key determinant of the rate that a new technology is adopted. In our model, learning at the plant level is represented by shocks to the plant's specific productivity. We argue here that data on plant size can uncover these shocks. We first document that plants grow as they age while their average productivity remains roughly constant. We then discuss how theory links plant-specific productivity to plant size and contrast our approach to the standard approach taken in the literature on learning. Finally, we discuss how we implement our method of measuring learning.

A. The Data on Plant Age, Size, and Average Productivity

Here we provide evidence on the two observations from the U.S. micro data on plants that motivate our model of learning.

The first observation is that plants grow substantially as they age. There is both cross-section and panel evidence on this point. In terms of the cross-section evidence, we look at the relative number of people employed by plants as they age. In Figure 4, we plot employment shares by plant cohorts from the 1988 panel of the U.S. Census Bureau's Longitudinal Research Database (LRD).[3] In the data, the employment share of plants rises at least for the first 20 years of a plant's life. Note that the employment share of a cohort of plants of age 20 is more than seven times that of the cohort of brand new plants.

In terms of the panel evidence, we present data from Jensen, McGuckin, and Stiroh (2001) to show that the employment share of a cohort of plants starts small and grows with age relative to the total employment in manufacturing. Jensen, McGuckin, and Stiroh use the LRD to construct a balanced panel of manufacturing plants for 1963–92. These plants are organized into cohorts that enter in a given five-year period. Plants are included in the sample only if they survive for the entire period of the panel. In Figure 5, we plot the ratio of average hours worked relative to the industry average for each of five such cohorts of plants. Clearly, the employment share in plants starts relatively small and grows steadily with age. Moreover, this pattern is evident for each of the five cohorts of plants.

The second observation behind our model of plant learning is that, after industry effects are controlled for, the average productivity of capital and labor does not vary systematically with plant age or size. This is documented by Bartelsman and Dhrymes (1998) and Jensen, McGuckin, and Stiroh (2001). Bartelsman and Dhrymes (1998) study a sample of manufacturing plants drawn from the U.S. Census Bureau's LRD. For this sample, they

report a geometric weighted-average of capital and labor productivity:

$$\left(\frac{y_{it}}{l_{it}}\right)^{1-\alpha}\left(\frac{y_{it}}{k_{it}}\right)^{\alpha}.$$

They calculate this weighted-average by age category and size decile as measured by the average level of employment over the period 1972–86, where the weights are obtained from a regression of outputs on inputs. We graph their values for this measure by age categories and size deciles in Figure 6. While Bartelsman and Dhrymes (1998) find variations in this measure across individual plants, Figure 6 shows no systematic relationship between their measure of average productivity and either age or size. Likewise, Jensen, McGuckin, and Stiroh (2001) report that after about 5–10 years, all cohorts of surviving plants have similar productivity levels.

B. The Link Between Plant Size and Plant-Specific Productivity

Our model implies a tight link between plant size and plant-specific productivity.

To see this as simply as possible, consider a stripped-down version of our model in which the output in each plant $i = 1, \ldots, N$ in period t is given by

$$(9) \qquad y_{it} = z_t A_{it}^{1-\nu} l_{it}^{\nu},$$

where a plant's productivity A_{it} may depend on its age. Profit maximization implies that plants equate their marginal products; thus, in equilibrium,

$$(10) \qquad \frac{l_{it}}{l_t} = \frac{A_{it}}{A_t},$$

where $l_t = \sum_i l_{it}$ is aggregate employment and $A_t = \sum_i A_{it}$ is the aggregate of specific productivities of all plants. Thus, employment in plant i relative to aggregate employment is equal to the specific productivity of plant i relative to aggregate specific productivity.

In our model with physical capital, a similar result holds for individual plants and for cohorts of plants. To establish this result, consider the static problem of allocating a given amount of capital and labor across plants at a point in time. For a given distribution λ_t of organization capital, it is convenient to define

$$(11) \quad n_t(A) = \frac{A}{\bar{A}_t}$$

as the *size* of a plant of type (A, s) in period t, where

$$(12) \quad \bar{A}_t = \sum_s \int_A A\lambda_t(dA, s)$$

is the aggregate of the specific productivities across all plants. The variable $n_t(A)$ measures the size of the plant in terms of its capital or labor or output, in that the equilibrium allocations are

$$(13) \quad k_t(A) = n_t(A)k_t, \; l_t(A) = n_t(A)l_t, \text{ and } y_t(A) = n_t(A)y_t.$$

To see this, note that since the production function F is homogeneous of degree 1, all operating plants in this economy use physical capital and labor in the same proportions. The proportions are those that satisfy the resource constraints for capital and labor.

Now define the aggregate of the specific productivities of a cohort of plants of age s as $\bar{A}_{t,s} = \int_A A\lambda_t(dA, s)$. Note from (11) that $\bar{A}_{t,s}/\bar{A}_t = \int_A n_t(A)\lambda_t(dA, s)$. Using (13), we then have this:

Proposition 2. *The aggregate of specific productivities of plants of age s relative to that of all plants is equal to the share of total employment in those plants; that is, $\bar{A}_{t,s}/\bar{A}_t = l_{t,s}$, where*

$$(14) \quad l_{t,s} = \int_A \frac{l_t(A)}{l_t}\lambda_t(dA, s).$$

Given this proposition, we can use the data on employment shares by cohorts, $l_{t,s}$, to infer the pattern of learning, as measured by $\bar{A}_{t,s}/\bar{A}_t$. From the perspective of our model, the LRD data in Figure 4 imply that the aggregate of specific productivities of a cohort of plants grows faster than aggregate productivity for at least 20 years. Since in the data the employment share of a cohort of plants of age 20 is more than seven times that of brand new plants, our model implies that plants which survive 20 years are, at that age, much more productive not only than they were when they were first built, but also than brand new plants in that 20th year. Thus, for a relatively long period of time, the ongoing innovations that occur within an operating plant on average are much larger than the innovations from the frontier technology. In this sense, 20-year-old plants are much less outdated than their contemporary brand new plants.

It is important to realize that the link between the employment shares $l_{t,s}$ and relative productivities $\bar{A}_{t,s}/\bar{A}_t$ established in Proposition 2 is independent of any changes or trends in overall employment l_t. Hence, the fact that in the data there are trends in manufacturing employment—increasing in the first part of the century and decreasing in the second part—has no bearing on our method of inferring learning from employment shares.

We have used employment shares to infer the amount of learning that plants experience as they age. We can also use them to infer the speed of diffusion of new technologies. In (8) we have defined diffusion in terms of output. Clearly, from (13), this diffusion is equal to the fraction of labor employed in plants using these technologies, so that

$$(15) \qquad D_{t,t+k} = \sum_{s=0}^{k} l_{t,s}.$$

C. A Contrast With the Literature on Learning

We have argued that theory implies that the plant-specific, or organization-specific, component of productivity can be uncovered from data on the relative size of organizations. Other researchers have not used these data to infer learning. Instead, Jovanovic and Nyarko (1995), Hornstein and Krusell (1996), Cooley, Greenwood, and Yorukoglu (1997), Greenwood and Yorukoglu (1997), and others rely on estimates of learning at the plant level made by Bahk and Gort (1993). We argue that these estimates are based on a conceptually flawed analogy to the early literature on learning for individuals performing specific tasks or groups of individuals performing a given number of tasks.

This early literature shows that in a wide variety of tasks, an individual's average productivity increases with the number of times that task is performed. Bahk and Gort's (1993) approach is based on the reasoning that because this is true for individuals, it should be true for groups of individuals, or organizations: average labor productivity in an organization should increase as the plant produces more output. Specifically, Bahk and Gort run a regression at the plant level of plant output on plant inputs and some measure of experience and interpret the coefficient on the experience variable as measuring the extent of learning. Unfortunately, this approach is valid only if the movement in plants' inputs is essentially unrelated to their specific productivity. Theory, however, predicts precisely the opposite, as the following example makes clear.

Consider running Bahk and Gort's regression in a simplified version of our model. In this simplified model, let output in plant i in period t be given by (9), so that relative employment in this plant is given by (10). Hence, taking logs of (9) and substituting for A_{it}

from (10) gives that, in equilibrium,

$$(16) \quad \log y_{it} = \log[z_t(A_t/l_t)^{1-\nu}] + \log l_{it}.$$

We can use (16) to calculate the coefficients in the following regression, of the form used by Bahk and Gort (1993):

$$\log y_{it} = \beta_{1t} + \beta_2 \log l_{it} + \beta_3 x_{it},$$

where x_{it} is some measure of the age or past production of the plant. This regression necessarily yields estimates of $\beta_3 = 0$ along with $\beta_{1t} = \log[z_t(A_t/l_t)^{1-\nu}]$ and $\beta_2 = 1$. Bahk and Gort's interpretation of $\beta_3 = 0$ would be that there is no learning at the plant level. This is true regardless of the amount that specific productivity A_{it} rises with age.

Notice also that in the equilibrium of the model, (16) implies that labor productivity is given by $y_{it}/l_{it} = z_t(A_t/l_t)^{1-\nu}$ and, hence, is constant across all plants, regardless of their specific productivity A_{it}. This observation points to the key difference between the implications for learning by an individual and that of an organization which can add variable factors. Individuals who learn increase their labor productivity. Organizations that learn grow by adding variable factors so as to keep their labor productivity constant (at least with Cobb-Douglas production). Hence, we argue that the key variable to look at to determine the amount of learning, or built-up knowledge or organization-specific capital, is not some measure of either capital or labor productivity but rather some measure of relative size.

D. Choosing the Learning Parameters

We choose the parameters governing the shocks to specific productivity so that the model matches data on the fraction of the labor force employed in plants of different age

groups, as well as data on job creation and job destruction in plants of different age groups, from the 1988 panel of the U.S. Census Bureau's LRD.[4] We choose the data from this panel because it has the most extensive breakdown of plants by age. We solve for the equilibrium of the model, including the exit decisions, and choose parameters so that the simulated moments from the model match those in the data. We need to solve for the equilibrium of the model in order to set the learning parameters because plants exit.

To get some intuition for how our procedure for choosing parameters works, it is useful to simply suppose, for a moment, that plants never exit. In such a case, there is a simple map between the underlying parameters of the learning process and data on employment shares and job creation and destruction of different age cohorts. We discuss this case in detail below because it is useful for building intuition for our procedure. We then describe the details of our procedure which deals with the relevant case in which plants do exit. In this case, we must solve for the entire equilibrium to choose the learning parameters.

The Case of No Exit

Suppose that plants never exit. In such a case, it is simple to choose the learning parameters since (11) implies that the size of a plant is directly proportional to its specific productivity. Hence, we can directly infer the dynamics of specific productivity from the dynamics of measured plant size.

Specifically, since the shocks to specific productivity ε_{is} for plant i from cohort s satisfy $A'_{is+1} = \varepsilon_{is} A_{is}$, (13) implies that

$$\frac{l_{is+1}}{l_{is}} \frac{l_s}{l_{s+1}} = \frac{A_{is+1}}{A_{is}} \frac{\bar{A}}{\bar{A}'} = \varepsilon_{is} \frac{\bar{A}}{\bar{A}'}, \tag{17}$$

where a prime denotes the next period's value. Hence, in an economy with no exit, the

distribution of the growth of relative employment for all plants in an age cohort s gives the whole distribution π_s of the shocks to specific productivity to that cohort up to the scalar $\bar{A}/\bar{A}'$, reflecting aggregate productivity growth.

To see how we can use data on employment shares, job creation, and job destruction in different age groups to set the learning parameters, suppose that the distribution of specific productivity shocks π_s is lognormal. Specifically, let $\log \varepsilon_{is} \sim N(\mu_s, \sigma_s)$, where μ_s and σ_s denote the age s–dependent mean and standard deviation of these shocks. Here we need only set the parameters μ_s and σ_s.

Data on employment shares are linked to these parameters as follows. Since shocks ε are independent over time, the growth rate from one period to the next of the fraction of the labor force employed in an age group is the mean of the distribution π_s up to a correction for growth in aggregate productivity. In particular, in a steady state,

$$(18) \qquad \frac{l_{s+1}}{l_s} = \frac{\exp(\mu_s + \sigma_s^2/2)}{1 + g_A},$$

where l_s is the steady-state employment of the cohort of age s, $1 + g_A$ is the steady-state growth rate of $\bar{A}_t$, and $\exp(\mu_s + \sigma_s^2/2)$ is the mean of π_s.

Data on job creation and destruction are linked to the learning parameters as follows. Job creation in a cohort of plants of age s is a measure of the dispersion of the distribution of specific productivity shocks π_s. In a steady state, any plant with a specific productivity shock greater than the growth of average productivity creates jobs. Hence, in an economy with no exit, the job creation rate for the cohort of plants of age s is

$$(19) \qquad C_s = \int_{\varepsilon > 1+g_A} \left(\frac{\varepsilon}{1 + g_A} - 1 \right) \pi_s(d\varepsilon),$$

which is some simple function, say, $H(\mu_s, \sigma_s)$, of the underlying parameters μ_s and σ_s. Job destruction is a similar measure of dispersion.

In sum, for any cohort s, we can use data on employment shares (l_{s+1}/l_s) and job creation rates C_s to back out the underlying learning parameters μ_s and σ_s.

The Case of Exit

The actual procedure we use to choose parameters is conceptually similar to that when plants never exit, but that procedure must be adjusted to deal with the fact that both in the model and in the data, plants do exit. This possibility of exit introduces a selection issue because the observed distribution of sizes is an endogenously truncated version of the distribution of plant-specific productivities. To deal with this endogenous selection issue, we need to solve the equilibrium of the model to solve for the exit decision. In practice, we choose the parameters μ_s and σ_s so that, in equilibrium, the model's predictions for employment shares, job creation, and job destruction of plants of different cohorts match the data. We explain this procedure in detail in Appendix B.[5]

In Figure 7, we display graphs which demonstrate how well our model with our chosen parameters reproduces the U.S. data. We plot actual employment and rates of job creation and destruction for U.S. manufacturing establishments in 1988, for plants of various ages, along with the comparable statistics generated by our model. The model with our chosen parameter values can reproduce these features of the data fairly well. For completeness, note that the implied statistics for the overall job creation and destruction rates expressed as percentages of total employment are 8.3% and 8.4% for the data and 9.9% and 9.9% in the model. The differences between the overall job creation and destruction rates in the data

and the model are not large compared to the fluctuation in these rates in annual data for the period 1972–93. The standard deviations of the job creation and job destruction rates over that period are 2.0 and 2.7.

In Figure 8, we display measures of the learning parameters themselves. We plot the mean shock to the size of the cohort of plants of age s relative to the industry average size, $\mu_s - g_A$. Here μ_s is the mean of the log of specific productivity shocks and g_A is the growth of average productivity. For the first 20 years, μ_s is greater than g_A, so that plants in cohort s grow relative to the industry average, while after that, they shrink. We also plot the standard deviations of the shocks to the log of specific productivity, which also equals the standard deviation of the shocks to the log of plant size. As the figure shows, these are roughly constant with age.

4. Choosing the Macro Parameters

Our choice of macro parameters is standard. The growth rate of output per hour g, the physical capital share $\nu\theta$, and the depreciation rate δ are chosen to reproduce data on the U.S. manufacturing sector. We set $g = 3.3\%$ to match the growth of manufacturing output per hour in 1949–69, as reported in Figure 2. We use data for 1959–99 obtained from the U.S. Department of Commerce's national income and product accounts to set $\nu\theta = 18.4\%$ and $\delta = 7.7\%$, based on methodology described by Atkeson and Kehoe (2002). We set $\beta = 0.977$, so that the steady-state interest rate i defined by $1 + i = (1 + g)/\beta$ is 5.7%.

Consider next the growth of the Solow residual. The steady-state growth rate of output per worker, $1 + g$, is related to the growth of the Solow residual by $(1 + g)^{1-\nu\theta}$, which can be decomposed as $(1 + g)^{1-\nu\theta} = (1 + g_z)(1 + g_\tau)^{1-\nu}$, where g_z and g_τ are the growth rates of

disembodied technology z_t and the frontier technology τ_t. Given our choices of $g = 3.3\%$ and $\nu\theta = 18.4\%$, using $(1+g)^{1-\nu\theta} = 1.027$ implies that the growth of the Solow residual is 2.7%. Since we calibrate our model to reproduce observations on plant size, the steady state is not affected by the decomposition of the Solow residual into these components.

Finally, consider the span of control parameter ν. Hundreds of studies have estimated production functions with micro data. These analyses incorporate a wide variety of assumptions about the form of the production technology and draw on cross-section, panel, and time series data from virtually every industry and developed country. Douglas (1948) and Walters (1963) survey many studies. Atkeson, Khan, and Ohanian (1996) review this literature and present evidence, in the context of a model like ours, that $\nu = 0.85$ is a reasonable value for this parameter. We use this value in our main analysis, but we also explore the sensitivity of our results to alternative values.

5. The Transition to a New Economy

Now we use our quantitative model to simulate the transition to a new economy with a permanently faster pace of technical change. We think of this simulation as a quantification of the hypotheses of Devine (1983), David (1990), David and Wright (1999), and others for the slow transition after the Second Industrial Revolution.

In the simulation, the increase in technical change is driven by faster growth in the frontier blueprints for new plants. We think of this faster growth of the frontier blueprints as capturing the faster pace of technical change embodied in plant design after the development of electric power. We find that the process of replacing old plants based on old blueprints with new plants is gradual. As a result, new embodied technologies diffuse slowly, and there

is a substantial delay before the increase in the pace of technical change leads to faster growth of output per hour. We then show that the rate of diffusion gradually accelerates from the old economy to the new one and argue that this implication is broadly consistent with the data. Finally, we use our model to isolate the economic forces that are essential in generating slow diffusion and transition. We repeat our transition experiment under alternative sets of assumptions to show the importance of our three key assumptions and the role of built-up knowledge.

A. The Transition Experiment

Consider an economy that is on a balanced growth path, with steady growth in the frontier blueprints causing output to grow 1.6% per year. This growth rate is the trend growth rate of output per hour in U.S. manufacturing for 1869–99 shown in Figure 2. We denote this initial growth rate of the frontier blueprints by g_τ^{old}. In our experiment, we suppose that at the beginning of the period labeled *1869,* agents learn that the growth rate of the frontier blueprints increases once and for all, so that on the new balanced growth path, output grows 3.3% per year (as in the U.S. data for 1949–69). We denote this new growth rate of the frontier blueprints by g_τ^{new}. We refer to these two balanced growth paths as the *old* and *new* economies, respectively.

In our experiment, we set the initial capital-output ratio and the distribution of organization capital across plants to be those from the balanced growth path of the old economy. In setting this initial distribution of organization capital, we assume that the distributions of the shocks to specific productivity are the same as those we estimated from the U.S. micro data for 1988.

Note, especially, two characteristics of this experiment. One is that the initial conditions for the experiment do not depend on the relative importance of embodied and economy-wide technical change in the old steady state. The other notable characteristic has to do with the fact that in the experiment, the shocks to specific productivities are held fixed while the growth rates of the frontier blueprints are varied. This amounts to assuming that the process of learning about any particular embodied technology does not depend on the rate at which new embodied technologies appear.

Consider now the diffusion of new technology during the transition to a new economy. In Figure 9, we show this diffusion in the model and in the data during 1869–1939. For the model, we graph the percentage of output produced in plants with blueprints dated 1869 and later. For the data, we graph the percentage of total horsepower in U.S. manufacturing establishments provided by electric motors. To make this comparison, we are assuming that in the data, plants that are driven by electric motors were built in and after 1869 and those driven by steam and water were built before 1869. The model predicts a slow pace of diffusion similar to that in the data. In the model, technologies dated 1869 and later take 45 years to diffuse to 50%; in the data, they take about 50 years.

Of course, the choice of initial dates in the data is somewhat arbitrary. To make a comparison that is not so dependent on initial dates, consider a statistic that is often used in the diffusion literature: the time it takes for diffusion to go from 5% to 50%. This time is roughly 20 years (1899–1919) in the data; it is 19 years in the model. Either way we measure it, therefore, the diffusion in the model is similar to that in the data.

In Figure 10, we show the model's implications for output per hour during the transition from 1869 to 1969, together with the actual data for those years. In the model as in

the data, the growth in output per hour gradually accelerates. Over the period 1869–99, the trend growth rate in output per hour is 1.6% in both the model and the data. Over the period 1899–1929, the trend growth rate in output per hour is 2.4% in the model and 2.6% in the data. In the model, the growth rate in output per hour reaches its new steady state of 3.3% by 1940.

B. Diffusion in the Old and New Economies

Our model implies a gradual acceleration in the speed of diffusion of new technologies: slow in the old economy, medium during the transition, and fast in the new economy. Here we argue that this implication is consistent with the data. Since we have just discussed the transition period, we only need to examine the old and new economies here.

In our model, in the old economy, a new technology takes about 68 years to diffuse to 50%; during the transition, this diffusion takes about 45 years; and in the new economy, it takes about 25 years.

In terms of the data, the slow diffusion of a new technology in the old economy is similar to the slow diffusion of steam power in the United States throughout the 1800s. The data in Figure 1 indicate that by 1869 steam power had diffused to just over 50%. If we assume that the diffusion of steam power started sometime between 1800 and 1810—as Atack, Bateman, and Weiss (1980) suggest—then these data indicate that steam power took roughly 60–70 years to diffuse to 50%. Moreover, as also shown in Figure 1, steam power took roughly another 20 years to diffuse from 50% to 80%. In our model, in the old economy, this same diffusion of a new embodied technology from 50% to 80% takes 19 years. Thus, the diffusion in our model's old economy is roughly consistent with U.S. data on the diffusion of

steam power in the 1800s.

In the new economy, new technologies take about 25 years to diffuse to 50%. We arrive at this diffusion rate by using (15) and the shares of employment for plants of different ages observed in the micro data from 1988 that we used in the calibration displayed in Figure 4.

C. Three Necessary Assumptions

Now we describe three assumptions without which our model would not generate the slow diffusion of new technology and the consequent slow transition to a new economy. All three of these assumptions are meant to capture the critical elements of the historians' hypotheses.

Our first necessary assumption is that most of the increase in the pace of technical change comes in the form of new technologies that are embodied in plants. To see that this assumption is necessary for the model to generate a slow transition, consider an alternative transition driven entirely by an increase in the growth rate of the economy-wide technology z and not by faster growth of the frontier blueprints. When the growth rate of the economy-wide technology increases, the production possibilities for all plants immediately increase with no loss of built-up knowledge. Suppose that at the beginning of 1869, agents learn that the growth rate of the economy-wide technology increases once and for all, so that on the new balanced growth path, output grows 3.3% per year. Here, the transition to a new economy is rapid: the trend growth rate of output during 1869–99 is 3.2%. Clearly, when the increase in the pace of technical change is not embodied in plants, the transition takes little or no time.

Our second necessary assumption is that improvements in the technology for new plants are ongoing. Clearly, for the model to generate a new economy with a higher growth

rate of output driven by embodied technical change, this technical change must be ongoing. If it were to stop, so would the growth due to embodied technical change.

Finally, our model's third necessary assumption is that after a new plant is built, it improves its technology during an extended period of learning. Without this assumption, diffusion of embodied technologies is very rapid. Specifically, if we set the age-dependent means of the shock to productivity to zero, then embodied technologies diffuse to 50% in just six years in the old economy and in three years in the new.

D. The Role of Built-Up Knowledge

For our model to generate an economic transition similar to that in the data, we find that we need one key ingredient in addition to the ingredients contained in the three assumptions discussed by the historians. The key ingredient is that there must be a large stock of *built-up knowledge* in the old economy. Here we discuss how to measure this stock of built-up knowledge, and we conduct two transition experiments that highlight its role.

A measure of the stock of built-up knowledge relative to the frontier blueprints is $(\bar{A}_t/\tau_t)^{1-\nu}$. Note that $\bar{A}_t/\tau_t$ is the average of the specific productivity across plants relative to the frontier blueprints available to new plants. The exponent $1-\nu$ expresses this ratio in units of the Solow residual of a standard growth model. In our transition experiment, this ratio is 1.24 in the new economy and 2.21 in the old. Thus, built-up knowledge is 78% higher in the old economy than in the new. The large stock of built-up knowledge in the old economy is what slows the transition to the new economy. As the transition begins, managers are reluctant to close existing plants and lose this knowledge for what, initially, is only a marginally superior technology.

We demonstrate the importance of this built-up knowledge for the speed of transition as follows. Suppose, counterfactually, that existing plants in 1869 have the stock of built-up knowledge corresponding to an economy with rapid growth of the frontier blueprints. Specifically, consider a transition in which the initial distribution of plant-specific productivities is the same as that in the new economy.

We can show analytically that if we also set the initial physical capital to output ratio equal to its steady-state value in the new economy, then there is no transition: the economy immediately enters a new steady state in which it grows at 3.3% per year, and new technologies diffuse to 50% in 25 years. If we solve the model with the initial capital-output ratio set to its steady-state value in the old economy, then the transition is essentially the same: the trend growth rate of this economy during 1869–99 is 3.2%, and technologies dated 1869 and later diffuse to 50% in only 27 years. Thus, the transition to a new economy occurs much faster when there is no large stock of built-up knowledge about old technologies.[6]

We do not have direct evidence on the stock of built-up knowledge in the old U.S. economy for the period we are studying. We do have some indirect evidence, however. In the context of our model, new technologies diffuse slowly when there is a large stock of built-up knowledge about existing technologies in plants. Given our model, the observation that steam power took roughly from 60 to 70 years to diffuse to 50% leads us to conclude that manufacturers had a substantial stock of built-up knowledge about existing technologies at the start of the Second Industrial Revolution. To quantify this initial stock of knowledge, we use our model to compute the distribution of organization capital that results when the frontier blueprints grow at a slow pace but the learning process, parameterized by the age-dependent distributions of shocks to specific productivity, is the same as in the new economy.

As demonstrated above, this assumption is key to generating a slow transition.

In the model, the stock of built-up knowledge in the old economy depends on the span of control parameter ν. As we increase that parameter, the stock of built-up knowledge in the old economy increases and the transition to the new economy slows. For example, consider our transition experiment with $\nu = 0.90$ as opposed to its baseline value of $\nu = 0.85$. In this experiment, the trend growth of output per hour in the model is 1.6% for the 1869–99 period and only 1.8% for the 1899–1929 period, and technologies dated 1869 and later take 57 years to diffuse to 50%. This transition is slower because with the higher value of ν, the old economy has 110% more built-up knowledge than the new economy. When $\nu = 0.85$, the comparable number is only 78%.

6. Parallels With the Recent Productivity Paradox

David (1990) argues that the transition to a new economy after the Second Industrial Revolution serves as a useful historical parallel for understanding the recent seeming paradox of rapid technical change in information technologies accompanied by relatively slow growth in productivity. We argue that while this comparison may be useful qualitatively, it is less useful quantitatively. We conclude that while the forces in our model can quantitatively account for the slow transition after the Second Industrial Revolution, some other forces must be at the heart of a quantitative explanation of the slow transition after the Information Technology Revolution.

In our model, the extent of built-up knowledge in existing organizations is smaller the faster is the pace of technical change. Since the pace of technical change was relatively fast even before the recent Information Technology Revolution began, our model implies that the

initial stock of built-up knowledge before this revolution is relatively small. Thus, the model counterfactually predicts that the speed of transition to a new economy should be relatively fast, not slow as it is in the data.

To make this concrete, consider another transition experiment. Suppose, as before, that the pace of technical change increases so that the steady-state growth rate increases 1.7 percentage points. But instead of starting with a relatively slow growth rate of 1.6% in the old economy, the economy starts with a relatively high growth rate of 3.3%. Suppose that in some period, agents learn that the growth of frontier blueprints has increased once-and-for-all, so that the economy grows 5% per year on the new balanced growth path. In this experiment, the trend growth of output per hour is 4.1% for the first 30 years and 5.0% for the next 30 years. In this transition, new technologies diffuse to 50% in only 14 years. Here the ratio of built-up knowledge in the old economy with 3.3% growth is only 20% higher than that in the new economy with 5% growth. Clearly, the transition to a new economy occurs significantly faster in this experiment than in our original experiment. Since technical change before the Information Technology Revolution was relatively fast, this experiment suggests that models like ours will predict a relatively fast transition to a new economy after this revolution, which is, again, counterfactual.

More generally, we can demonstrate the importance of the steady-state growth rates in the old economy for the speed of transition to a new economy by a simple experiment. We vary the initial growth rate in the old economy, but hold fixed the assumption that growth in the new economy is 1.7 percentage points higher than in the old. We show the results of this experiment in Figure 11. As the growth rate in the old economy increases, the stock of built-up knowledge in the old economy falls relative to that in the new economy, and the

speed of transition and diffusion both increase.

7. Extensions

Here we have used a very simple model that captures the essence of historians' hypotheses about the Second Industrial Revolution. Several features of our work perhaps deserve further remarks.

In the model, for example, we have abstracted from the possibility that agents could undertake investment to accelerate learning at the plant level. Writing down a theoretical model that allows for this possibility would be easy. Currently, though, it is unclear how to use independent evidence to discipline the parameters governing that part of the model.

In the model, we have also abstracted from sectors other than manufacturing. Theoretically, extending our model to include services and agriculture as well would be simple. But we don't think that such an extension would have a large impact on our results regarding the diffusion of new technologies embodied in manufacturing plants and the corresponding growth of average productivity of labor in manufacturing. If we added more sectors and conducted the same transition experiment, we would expect the interest rate effects emanating from the increased technical change in manufacturing to be diluted somewhat. In our current analysis, these effects are already quantitatively small.

In the model, we have also abstracted from fixed effects on plant size that are clearly present in the data. It would easy to include them by allowing the initial draw of plant productivity to differ across plants either within or across industries. From Proposition 2, it is clear that such fixed effects would have no impact on our measurements of learning for cohorts of plants.

In the model, we have included uncertainty about the evolution of plant-specific productivity. We have done so because the distribution of shocks to specific productivities determines the distribution of organization capital across plants and, hence, plays an important role in shaping the dynamics of transition. We could have ignored the substantial heterogeneity in these shocks readily apparent in the data on job creation and job destruction and instead assumed that these shocks were deterministic. We think that such an incomplete model would have misleading implications for the transition dynamics.

A final concern might be that the data have secular trends in the fraction of the labor force devoted to manufacturing. Because returns are diminishing at the plant level, the main way such trends would have an impact on the model's average labor productivity in manufacturing would be by changing the average number of workers per plant. We don't think such trends would play a quantitatively important role in accounting for the trends in labor productivity observed in the data.

8. Relation to the Literature

Our work here is related to several strands of literature.

The process of industry evolution and learning at the plant level in our model is related to that in the models of Jovanovic (1982), Hopenhayn (1992), Hopenhayn and Rogerson (1993), and Campbell (1998).

At the general level, the technology diffusion in our model is related to the theoretical literature on diffusion, including the work of Chari and Hopenhayn (1991) and Jovanovic and MacDonald (1994). An important feature that our model shares with the Chari-Hopenhayn model is that investment in old technologies continues even after the invention of a new tech-

nology. More recent theoretical work on diffusion after major technical change includes that of Aghion and Howitt (1998), Helpman and Trajtenberg (1998), and Laitner and Stolyarov (forthcoming). In this literature, the diffusion of a major new technology is constrained by the need to develop complementary inputs for that technology. In contrast, the point here is to investigate the hypotheses put forward by David and Wright (1999), that the new technologies stemming from electric power were embodied in plants and that the diffusion of these technologies was delayed by the reluctance of agents to abandon the knowledge they had built up in plants embodying older technologies. With the exception of Laitner and Stolyarov (forthcoming), a key difference between our work and this theoretical work on diffusion is that ours is quantitative. Moreover, most of these models are sufficiently abstract that attempting to use data to apply quantitative discipline to evaluate their implications is a daunting task.

More closely related to our work is the applied work of Hornstein and Krusell (1996) and Greenwood and Yorukoglu (1997). They are primarily interested in the post-1974 productivity slowdown. In their work, however, technologies are embodied in machines, not plants. The diffusion of new technologies embodied in new machines is slowed because agents must learn to use new machines efficiently. Both of these studies use the estimates of Bahk and Gort (1993) regarding learning at the plant level to calibrate the parameters in their models governing the learning process for new machines. As we stressed above, the Bahk and Gort estimates are conceptually flawed. Moreover, these models' predictions for the size-age profile of plants can easily be shown to be grossly at variance with the observations in the data. Finally, as all of these researchers note, their numerical simulations are intended to be only illustrative.[7]

Some related work has also been done on the skill premium and technology choice during a transition by Caselli (1999) and Krusell, Ohanian, Ríos-Rull, and Violante (2000).

Finally, literally hundreds of empirical studies examine the diffusion of individual products and processes that are not embodied in plant design. (For an interesting recent example studying the diffusion of tractors, see Manuelli and Seshadri 2003.) While these studies find a wide array of diffusion times, they are not germane to our work. Moreover, since we focus on technologies embodied in plant design, our model is not intended to be applicable to these other products and processes.

9. Conclusion

Here we have asked, can a model which incorporates quantitative measures of the learning process that match the U.S. micro data on the life cycle of plants account for the slow diffusion and slow transition in U.S. manufacturing after the Second Industrial Revolution? We have found that it can. This finding requires that much of the increase in the pace of technical change after the revolution come from new embodied technologies, that this increase be ongoing, and that there be substantial learning about these technologies at the plant level. In addition, we have discovered that for the model to generate slow diffusion and slow transition, manufacturers must start with a large stock of built-up knowledge about the existing technologies embodied in their plants. On the basis of this result, we have argued that the U.S. experience after the Second Industrial Revolution may not be quantitatively useful as a parallel with the U.S. experience after the Information Technology Revolution.

Of course, before any definitive analysis of the impact of the Information Technology Revolution can be fleshed out in a quantitative model, at least three issues must be addressed:

Where are the new technologies embodied? How long is the period of learning after they are adopted? And how much built-up knowledge do existing organizations have with their current technologies? With regard to information technologies, none of these questions are easy to answer.

In our model of the Second Industrial Revolution, we have followed the historical literature in assuming that the designs of new manufacturing plants embody new technologies. This assumption does not seem to be immediately applicable to the Information Technology Revolution, since where the information technologies are embodied is not clear. Some evidence suggests, in fact, that organizations can use these technologies efficiently only after the organizations have been substantially restructured, so the new technologies might be somehow embodied in the structure of the organization. (See Brynjolfsson and Hitt 2000.) Perhaps our model could be adapted to analyze the Information Technology Revolution, but the unit of analysis would probably have to shift to some level of organization other than plants.

In our model, we used data on the birth, growth, and death of plants to draw inferences about the speed of learning. It is not clear what corresponding data would help quantify the speed of learning during the Information Technology Revolution. Clearly, such data are critical to evaluating the impact of these technologies. One might argue that the relevant unit of analysis for the Information Technology Revolution should be firms rather than plants. In that regard, note that the data on the market shares of cohorts of firms are quite different from the corresponding data for plants. In particular, the market share for a cohort of firms declines rather than rises over time. (See the survey of Caves 1998.) In the context of our model, these data imply that the diffusion of new technologies embodied in firms is fast,

rather than slow. Hence, this avenue does not seem promising for accounting for the slow transition after the Information Technology Revolution.

Notes

[1]Note that we do not assume that all new technologies are embodied in new plants. In fact, our model is consistent with the idea that most of the ongoing technical change in the economy is driven by either disembodied technical change or technical change embodied in other factors, such as capital goods or labor.

[2]Since our model has a fixed number of managers and each manager can either start a new plant or operate an existing plant, our assumptions imply that on a balanced growth path, the number of plants is fixed. An alternative assumption, pursued by Hopenhayn and Rogerson (1993), is that instead what is fixed is the cost in terms of consumption goods of starting a new plant. In that alternative model, the number of plants grows over time.

We have chosen our specification because it seems to be a good approximation to the data. Sands (1961) reports that over the period 1904–47, the number of manufacturing plants in the United States grew only 0.5% per year while output per manufacturing establishment grew nearly 3.0% per year. Clearly, most of the growth of output in this period came from more output from each plant and only a small part from an increase in the number of plants.

[3]Here and throughout this study, our microeconomic data are taken from the U.S. Census Bureau's Longitudinal Research Database (LRD) on manufacturing plants, a data set available at http://www.census.gov. This data set is described in Davis, Haltiwanger, and Schuh (1996). We use data on employment, job creation, and job destruction from the 1988 panel of the LRD which we obtained from John Haltiwanger's Web site, http://www.bsos.umd.edu/econ/haltiwanger.

[4]Davis, Haltiwanger, and Schuh (1996) define the following statistics. *Job creation* in

a plant i in year t is $l_{it} - l_{it-1}$ if $l_{it} \geq l_{it-1}$ and zero otherwise, where l_{it} is the labor force in plant i in year t. *Job destruction* in a plant i in year t is $l_{it-1} - l_{it}$ if $l_{it} \leq l_{it-1}$ and zero otherwise.

[5]Notice that to make our model simple to exposit we have identified each plant as having one fixed factor of some given scale that is the same across plants. We have called this fixed factor a *manager*. We have assumed that this fixed factor gets a share $1 - v$ of value-added. Our current procedure also covers a more sophisticated case in which the level of the fixed factor varies across plants but the factor still receives a share of $1 - v$.

In particular, one way to deal with the observation that plants in some industries are systematically larger than plants in other industries is to assume that the scale of the fixed factor in some industries is larger than the scale in others. For example, the scale of the fixed factor may be ten units in an automobile plant and just one unit in a shoe plant. In this example, ten shoe plants would function like one automobile plant.

Notice that our current calibration procedure, as it stands, is consistent with allowing for such heterogenous level effects in the fixed factor. The reason is that in calibrating the learning process, we use only shares of employments by cohorts of plants and job creation and destruction rates. In our example, when faced with the same shocks, the ten shoe plants would produce identical observations on these variables as the one automobile plant. Thus, the key identifying assumptions in our procedure are two: the sum of the fixed factors in any given cohort of plants is constant over time, but can differ widely across plants within a cohort or across cohorts in different time periods; and the share of the fixed factor is constant.

[6]This transition experiment can be interpreted as the transition that occurs in an alternative model in which both learning and the growth of the frontier blueprints are slower

in the old economy than in the transition to the new economy. Specifically, assume that in the old economy, the means of the log of plant-specific productivity shocks are lower than those in the new economy by $\log(1 + g_\tau^{new}) - \log(1 + g_\tau^{old})$. In this alternative model, as the growth rate of the frontier blueprints increases, the rate of learning increases, so that the new economy has the same size-age distribution as the old economy.

[7]For example, Greenwood and Yorukoglu (1997, p. 67) say explicitly that their "model was not tuned, by choice of parameter values, to be in harmony with any particular features of the U.S. data."

Appendix A: Links Between Historical Analyses and the Model

Our model has three key assumptions: (1) new plants embody new technologies; (2) improvements in the technology for new plants are ongoing; and (3) new plants require an extended period of learning in order to use their technology efficiently. Motivating these assumptions are three features of the data discussed by historians in relation to the development and adoption of electricity that began around 1860: (1) the change to electric power from steam and water power led to major changes in factory design and machine organization that went hand in hand with electrification; (2) the process of improving efficiency through changes in factory design continued for decades; and (3) for each new factory design, the process of learning how best to use the new design took an extended period of time.

New Plants Embody New Technologies

Our assumption that new plants embody new technologies in their plant design is motived by the analysis of Devine (1983) and David (1990, 1991), who argue that the adoption of the modern technology of electricity required a complete redesign of the manufacturing plant. In steam- and water-driven plants, power was distributed mechanically throughout the factory by a series of shafts and belts called a *direct-drive system*. In modern electric plants, power is distributed as electricity through wires to individual motors in what is called a *unit-drive system*.

Devine (1983, pp. 350, 352) describes the direct-drive system used in steam- and water-driven plants this way:

> Until late in the nineteenth century, production machines were connected by a direct mechanical link to the power sources that drove them. In most factories,

> a single centrally located prime mover, such as a water wheel or steam engine, turned iron or steel "line shafts" via pulleys and leather belts. These line shafts—usually 3 inches in diameter—were suspended from the ceiling and extended the entire length of each floor of a factory, sometimes even continuing outside to deliver power to another building. Power was distributed between floors of large plants by belts running through holes in the ceiling . . . . The line shafts turned, via pulleys and belts, "countershafts"—shorter ceiling-mounted shafts parallel to the line shafts. Production machinery was belted to the countershafts and was arranged, of necessity, in rows parallel to the line shafts . . . . The entire network of line shafts and countershafts rotated continuously—from the time the steam engine was started up in the morning until it was shut down at night—no matter how many machines were actually being used. If a line shaft or the steam engine broke down, production ceased in a whole room of machines or even in the entire factory until repairs were made.

Panel A of Figure 12 (from Devine 1983) illustrates this direct-drive system of power.

The direct-drive system of power gradually evolved into the unit-drive system used in modern electricity-driven plants. In the unit-drive system, illustrated in panel D of Figure 12, each machine is driven by its own electric motor, and power for that motor is delivered through power lines from some potentially far-off electric utility plant. Panels B and C of Figure 12 show two short-lived intermediate stages in this evolution, referred to as *electric line shaft drive* and *electric group drive*. Panel A of Figure 13 (also from Devine 1983) shows the evolution of these methods of driving machinery with electricity in manufacturing. Panel

B shows the key developments of electric motors and power generation.

Redesigning and Learning Take Time

Our other two assumptions—that technical change in plant design is ongoing and that the process of learning how best to use each new factory design takes an extended period of time—are motivated by the work of many historians.

As Devine (1983) describes, with the unit-drive system, plants could be designed and machinery in the plant could be arranged so as to handle materials according to the natural sequence of manufacturing operation, rather than according to physical placement of shafts, as required by the direct-drive system. Moreover, once the shafts in the direct-drive system became unnecessary, plants could also be designed with improved ventilation, illumination, and cleanliness and to accommodate overhead electric cranes, which were thought to revolutionize materials-handling.

Devine (1990) and Sonenblum (1990) document that the adoption of the unit-drive system was only the first of a series of ongoing advances in the production process and modes of organization of factories that depended on the use of electricity. Moreover, these historians argue that each advance in factory design required an extended period of learning how to best use the new design.

Sonenblum (1990) documents three stages of factory design evolution. In the first stage (which he says occurred in 1899–1920), new factory design evolved from the traditional direct drive through the intermediate stages of the electric line shaft drive and the electric group drive to the electric unit drive. In the second stage (1920–48), the attention in new factory design shifted to modifying the factory layout in order to accelerate the throughput

of materials. Machines were arranged so that materials moved smoothly from one operation to the next, and assembly lines became more common. In the third stage (1948–85), new factories were designed to optimally use new servomechanisms that automatically controlled machine actions and numerically controlled machines. While these factory systems were developed in the 1940s and 1950s, they began to spread into many manufacturing plants only in the 1960s. These systems allowed factories to rely less on large, inflexible assembly lines and to produce nonstandard products in small batches. Moos (1957) and Slesinger (1958) also discuss the changes in plant design driven by the development of automatically controlled machines and the learning required to take advantage of such plants. As Devine (1990) discusses, in the 1980s, the evolution of factory design evolved to accommodate new methods of computer materials-handling and computer-integrated manufacturing in which a computer controls whole groups of machines. Figure 14 (from Devine 1990) gives a brief chronology of this ongoing change in the organization of production that required new factory designs.

Chandler (1992, p. 84) discusses the type of built-up organizational capabilities that resulted from firms learning to efficiently use the technologies developed in the Second Industrial Revolution. He argues that the learned capabilities that resulted from solving problems of scaling up the processes of production manifest themselves in firms' production and distribution facilities. These learned capabilities were developed through trial and error, feedback, and evaluation and were organization-specific.

Appendix B: Choosing Micro Parameters

In our calibration, we assume that the economy is on a balanced growth path. Here we show that along a balanced growth path, the Bellman equations characterizing the entry and exit decisions of plants, (4) and (5), can be recast with plant size n replacing specific productivity as a state variable. In the recast Bellman equation, the only preference or technology parameters that enter are the discount factor β and the age-dependent distributions of shocks to size. This result allows us to calibrate the shocks to specific productivity A directly in terms of shocks to plant size n and, moreover, to do so in a way independent of the rest of the parameters in the model. After presenting this result, we discuss our choice of these parameters in greater detail.

Recasting the Bellman Equations

Consider, first, the balanced growth path. To ensure that our model has such a path, we assume that $F(k,l)$ has the Cobb-Douglas form $k^{\theta}l^{1-\theta}$. We define a *balanced growth path* in this economy as an equilibrium in which the following conditions hold: The quality of the frontier blueprint τ_t and the productivity $\bar{A}_t$ grow at a constant rate $1+g_{\tau}$; the economy-wide level of technology z_t grows at a constant rate $1+g_z$; aggregate variables y_t, c_t, k_t, w_t, and w_{mt} grow at a rate $1+g$, where $1+g = [(1+g_z)(1+g_{\tau})^{1-\nu}]^{1/(1-\nu\theta)}$; and variables ϕ_t, V_t^0, and r_t are constant. Moreover, the distribution of plants by age and size is constant over time, and the value of a plant of a given age and size grows at the rate $1+g$, so that $\mu_{t+1}(A',s) = \mu_t\big(A,s\big)$ and $V_{t+1}(A',s) = (1+g)V_t\big(A,s\big)$ for all t, A, s with $A' = A(1+g_{\tau})$. Similar relations hold for λ_{t+1}, V_{t+1}^c, and d_{t+1}.

Along the balanced growth path, we can recast our state variables as (n,s) instead of

(A, s) as follows. Let $\{\rho_s\}$ be the cumulative distribution functions of $\eta = \epsilon/(1+g_\tau)(1+g_z)$ induced by $\{\pi_s\}$. We refer to $\{\rho_s\}$ as the steady-state distributions of shocks to plant size. Consider the Bellman equation

$$(20) \quad W(n,s) = \max\left[0, W^c(n,s)\right]$$

$$W^c(n,s) = n - \omega_m + \beta \int_\eta W(n\eta, s+1)\rho_{s+1}(d\eta),$$

where $\omega_m = \beta W(1,0)$.

Proposition 3. *The distribution of plants by size and age on the balanced growth path is determined by the solution to* (20).

Before we prove this proposition, note that the only parameter that enters (20) besides parameters governing the shocks to size is the discount factor β. Therefore, given β, we can choose the parameters of the distributions of the shocks to plant size to match data on the size and age distribution of plants independently of our choices for the other macro parameters.

Proof. We use the solution to (20) to construct a solution to (4) as follows. First, we use the function $W^c(n,s)$ to define an operating rule for plants contingent on state n and age s: plants with $W^c(n,s) \geq 0$ operate, and those with $W^c(n,s) < 0$ do not. We use this operating rule, together with the distributions of shocks to size $\{\rho_s\}$ and the normalization that new plants start with $n = 1$ and $s = 0$ to compute the implied distribution of n across plants of different ages s, which we denote by $\eta(n,s)$. To convert the state variable n to correspond to the size of the plant and the distribution $\eta(n,s)$ to correspond to the size-age distribution of plants, we normalize n so that the distribution of size across plants integrates to one.

To that end, let $\bar{n} = \sum_s \int_n \eta(dn, s)$. Now normalize periods on the balanced growth path so that $\tau_0 = 1$, and choose the units of economy-wide productivity so that $z_0 = 1$. Define $\bar{A}_0 = \bar{n}$, and solve for the corresponding values of y_0 and k_0 on the balanced growth path with growth rate g. Define the function $V_0(A, s) = W(n, s)y_0(1 - \nu)$, where $n = A/\bar{A}_0$. Then use the balanced growth relation, $V_{t+1}(A', s) = (1 + g)V_t\big(A, s\big)$ for all t, A, s with $A' = A(1 + g_\tau)$, to define the function $V_t(A, s)$. Define the functions $V_0^c(A, s)$ and $d_t(A)$ and the sequence w_{mt} similarly. Also, construct the sequence $\bar{A}_t$ using $\bar{A}_{t+1} = (1 + g_\tau)\bar{A}_t$. That these functions $V_t(A, s)$, $V_t^c(A, s)$, $d_t(A)$, and wages w_{mt} solve the original Bellman equation (4) along the steady-state path follows from the fact that the functions W and W^c defined above satisfy this second Bellman equation (20). That $\bar{A}_0 = \bar{n}$ is the initial aggregate of specific productivities across plants follows from the normalization $\tau_0 = 1$, since $\tau_0/\bar{A}_0 = 1/\bar{n}$ is the size of new plants. Q.E.D.

Quantifying Shocks

We use microeconomic data to quantify the shocks to plant size η. Note that since only the product $(1 + g_\tau)(1 + g_z)$ enters the definition of shocks to size η, the data on the size-age distribution of plants do not pin down the relative contribution to growth in the Solow residual of growth in the two types of technology: frontier and economy-wide. More generally, the balanced growth properties of this model do not depend on the way a given growth in the Solow residual g is split between g_τ and g_z. However, the transition paths turn out to depend crucially on whether the transition is driven by changes in g_τ or by changes in g_z.

We use observations from micro data on manufacturing plants in the United States

to choose the parameters affecting the shocks to size using a type of simulated method of moments. We parameterize the distributions of these shocks as follows. We assume that shocks to size have a lognormal distribution, so that $\log(\eta_s) \sim N(m_s, \sigma_s^2)$. We choose the means and standard deviations of these distributions to be smooth functions of the age s of a plant. In particular, we set $m_s = \gamma_1 + \gamma_2(\frac{S-s}{S})^2$ for $s \leq S$ and $m_s = \gamma_1$ otherwise and $\sigma_s = \gamma_3 + \gamma_4(\frac{S-s}{S})^2$ for $s \leq S$ and $\sigma_s = \gamma_3$ otherwise. With this parameterization, the shocks to size for plants of age S or older are drawn from a single distribution. Thus, shocks to size are parameterized by $\{\gamma_i\}_{i=1}^4$ and S. Note that the corresponding shocks to plant-specific productivity are lognormal with means $\hat{m}_s = m_s + \log(1 + g_\tau)$ and variances $\hat{\sigma}_s^2 = \sigma_s^2$.

We choose the parameters governing the shocks to size so that the model matches data on the fraction of the labor force employed in plants of different age groups, as well as data on job creation and job destruction in plants of different age groups, from the 1988 panel of the U.S. Census Bureau's LRD. We use the data from this panel because it has the most extensive breakdown of plants by age.

To produce the shocks' parameter values, we set the model's parameter $S = 150$ and choose the γ_i to minimize the sum of the squared errors between the model's and the data's statistics.

The parameters that generate these shocks are $S = 150$, $\gamma_1 = -0.1843, \gamma_2 = 0.2481$, $\gamma_3 = 0.1888$, and $\gamma_4 = 0.0005$. Recall that in Figure 8, we have plotted the means and standard deviations of our estimated shocks to the log of the size of plants, m_s and σ_s, that are implied by these parameters for $s = 0, \ldots, 25$.

References

Aghion, Philippe, and Howitt, Peter. 1998. *Endogenous growth theory.* Cambridge Mass.: MIT Press.

Atack, Jeremy; Bateman, Fred; and Weiss, Thomas. 1980. The regional diffusion and adoption of the steam engine in American manufacturing. *Journal of Economic History* 40 (June): 281–308.

Atkeson, Andrew, and Kehoe, Patrick J. 2002. Measuring organization capital. Working Paper 8722. National Bureau of Economic Research.

Atkeson, Andrew; Khan, Aubhik; and Ohanian, Lee. 1996. Are data on industry evolution and gross job turnover relevant for macroeconomics? *Carnegie-Rochester Conference Series on Public Policy* 44 (June): 216–50.

Bahk, Byong-Hyong, and Gort, Michael. 1993. Decomposing learning by doing in new plants. *Journal of Political Economy* 101 (August): 561–83.

Bartelsman, Eric J., and Dhrymes, Phoebus J. 1998. Productivity dynamics: U.S. manufacturing plants, 1972–1986. *Journal of Productivity Analysis* 9 (January): 5–34.

Brynjolfsson, Erik, and Hitt, Lorin M. 2000. Beyond computation: Information technology, organizational transformation and business performance. *Journal of Economic Perspectives* 14 (Fall): 23–48.

Campbell, Jeffrey R. 1998. Entry, exit, embodied technology, and business cycles. *Review of Economic Dynamics* 1 (April): 371–408.

Caselli, Francesco. 1999. Technological revolutions. *American Economic Review* 89 (March): 78–102.

Caves, Richard E. 1998. Industrial organization and new findings on the turnover and mobility of firms. *Journal of Economic Literature* 36 (December): 1947–82.

Chandler, Alfred D. 1992. Organizational capabilities and the economic history of the industrial enterprise. *Journal of Economic Perspectives* 6 (Summer): 79–100.

Chari, V.V., and Hopenhayn, Hugo. 1991. Vintage human capital, growth, and the diffusion of new technology. *Journal of Political Economy* 99 (December): 1142–65.

Cooley, Thomas F.; Greenwood, Jeremy; and Yorukoglu, Mehmet. 1997. The replacement problem. *Journal of Monetary Economics* 40 (December): 457–99.

David, Paul A. 1990. The dynamo and the computer: An historical perspective on the modern productivity paradox. *American Economic Review* 80 (May): 355–61.

_____. 1991. Computer and dynamo: The modern productivity paradox in a not-too-distant mirror. In *Technology and productivity: The challenge for economic policy,* pp. 315–45. Paris: Organisation for Economic Co-operation and Development.

David, Paul A., and Wright, Gavin. 1999. General purpose technologies and surges in productivity: Historical reflections on the future of the ICT revolution. Department of Economics Working Paper 99-026. Stanford University.

Davis, Steven J.; Haltiwanger, John C.; and Schuh, Scott. 1996. *Job creation and destruction.* Cambridge, Mass.: MIT Press.

Devine, Warren D., Jr. 1983. From shafts to wires: Historical perspective on electrification. *Journal of Economic History* 43 (June): 347–72.

_____. 1990. Electricity in information management: The evolution of electronic control. In *Electricity in the American economy: Agent of technological progress*, ed. Schurr, Sam H.; Burwell, Calvin C.; Devine, Warren D., Jr. and Sonenblum, Sidney.

Westport, Conn.: Greenwood Press.

Douglas, Paul H. 1948. Are there laws of production? *American Economic Review* 38 (March): 1–41.

Gordon, Robert J. 2000a. Does the "new economy" measure up to the great inventions of the past? *Journal of Economic Perspectives* 14 (Fall): 49–74.

_____. 2000b. Interpreting the "one big wave" in U.S. long-term productivity growth. In *Productivity, technology, and economic growth,* ed. Bart van Ark, Simon K. Kuipers, and Gerard H. Kuper, pp. 19–65. Boston: Kluwer.

Greenwood, Jeremy, and Yorukoglu, Mehmet. 1997. 1974. *Carnegie-Rochester Conference Series on Public Policy* 46 (June): 49–95.

Helpman, Elhanan, and Trajtenberg, Manuel. 1998. Diffusion of general purpose technologies. In *General purpose technologies and economic growth,* ed. Helpman, Elhanan, pp. 85–119. Cambridge, Mass.: MIT Press.

Hopenhayn, Hugo A. 1992. Entry, exit, and firm dynamics in long run equilibrium. *Econometrica* 50 (September): 1127–50.

Hopenhayn, Hugo, and Rogerson, Richard. 1993. Job turnover and policy evaluation: A general equilibrium analysis. *Journal of Political Economy* 101 (October): 915–38.

Hornstein, Andreas, and Krusell, Per. 1996. Can technology improvements cause productivity slowdowns? In *NBER Macroeconomics Annual 1996,* ed. Ben S. Bernanke and Julio J. Rotemberg, pp. 209–59. Cambridge, Mass.: MIT Press.

Jensen, J. Bradford; McGuckin, Robert H.; and Stiroh, Kevin J. 2001. The impact of vintage and survival on productivity: Evidence from cohorts of U.S. manufacturing plants. *Review of Economics and Statistics* 83 (May): 323–32.

Jovanovic, Boyan. 1982. Selection and the evolution of industry. *Econometrica* 50 (May): 649–70.

Jovanovic, Boyan, and MacDonald, Glenn M. 1994. Competitive diffusion. *Journal of Political Economy* 102 (February): 24–52.

Jovanovic, Boyan, and Nyarko, Yaw. 1995. A Bayesian learning model fitted to a variety of empirical learning curves. *Brookings Paper on Economic Activity,* Microeconomics, pp. 247–305.

Krusell, Per; Ohanian, Lee E.; Ríos-Rull, José-Víctor; and Violante, Giovanni L. 2000. Capital-skill complementarity and inequality: A macroeconomic analysis. *Econometrica* 68 (September): 1029–54.

Laitner, John, and Stolyarov, Dmitri. Forthcoming. Technological change and the stock market. *American Economic Review.*

Lucas, Robert E., Jr. 1978. On the size distribution of business firms. *Bell Journal of Economics* 9 (Autumn): 508–23.

Manuelli, Rodolfo, and Seshadri, Ananth. 2003. Frictionless technology diffusion: The case of tractors. Working Paper 9604. National Bureau of Economic Research.

Moos, S. 1957. The scope of automation. *Economic Journal* 67 (March): 26–39.

Rosenberg, Nathan. 1976. On technological expectations. *Economic Journal* 86 (September): 523–35.

Sands, Saul S. 1961. Changes in scale of production in United States manufacturing industry, 1904–1947. *Review of Economics and Statistics* 43 (November): 365–68.

Schurr, Sam H.; Burwell, Calvin C.; Devine, Warren D., Jr.; and Sonenblum, Sidney, ed. 1990. *Electricity in the American economy: Agent of technological progress.* Westport,

Conn.: Greenwood Press.

Schurr, Sam H.; Netschert, Bruce C.; Eliasberg, Vera; Lerner, Joseph; and Landsberg, Hans H. 1960. *Energy in the American economy, 1850–1975: An economic study of its history and prospects.* Baltimore: Johns Hopkins Press.

Slesinger, Reuben E. 1958. The pace of automation: An American view. *Journal of Industrial Economics* 6 (June): 241–61.

Sonenblum, Sidney. 1990. Electrification and productivity growth in manufacturing. In *Electricity in the American economy: Agent of technological progress*, ed. Schurr, Sam H.; Burwell, Calvin C.; Devine, Warren D., Jr.; and Sonenblum, Sidney. Westport, Conn.: Greenwood Press.

U.S. Department of Commerce, Bureau of Economic Analysis. 1973. *Long-term economic growth, 1860–1970.* Washington, D.C.: Bureau of Economic Analysis.

Walters, Alan A. 1963. Production and cost functions: An econometric survey. *Econometrica* 31 (January–April): 1–66.

Figures 1–2 **The Gradual Transition to a New U.S. Economy**

Figure 1 **The Slow Diffusion of Electric Power . . .**

Percentage of Total Horsepower From Three Sources of Mechanical Drive
in U.S. Manufacturing, 1869–1939

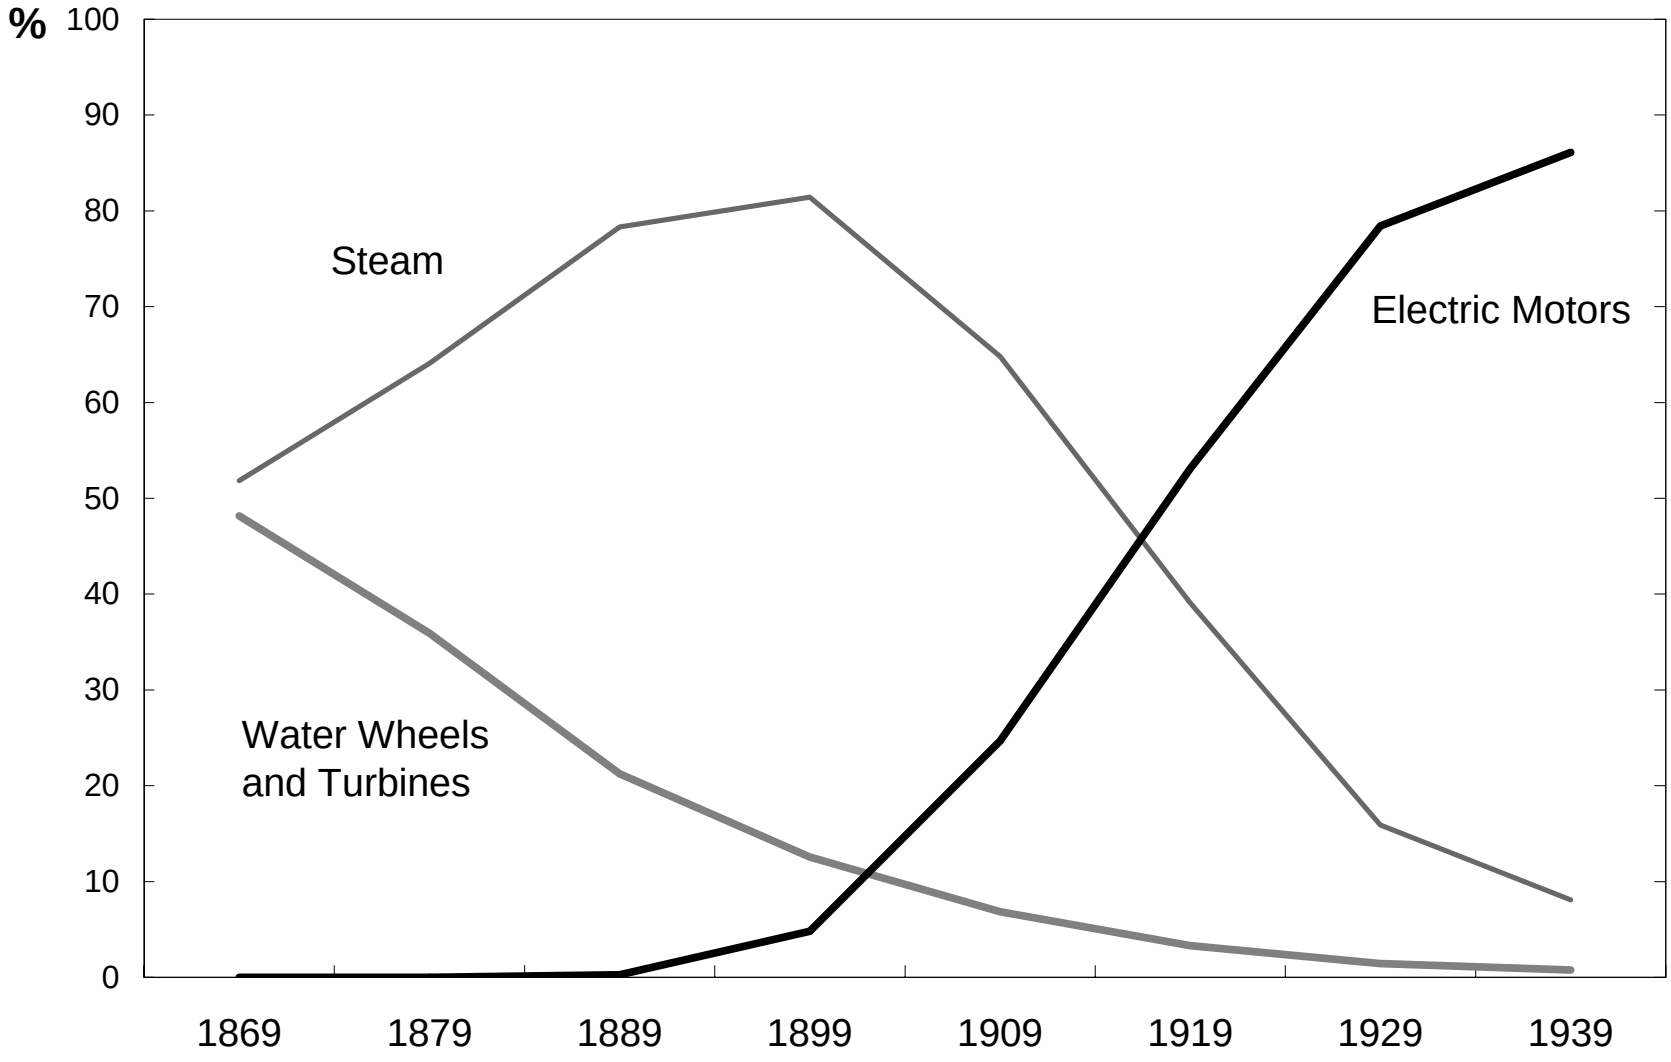

Source: Devine (1983, p. 351, Table 3)

Figure 2 **. . . And the Slow Increase in Productivity Growth**

Log of Output per Hour and Trend Growth in U.S. Manufacturing, 1869–1969

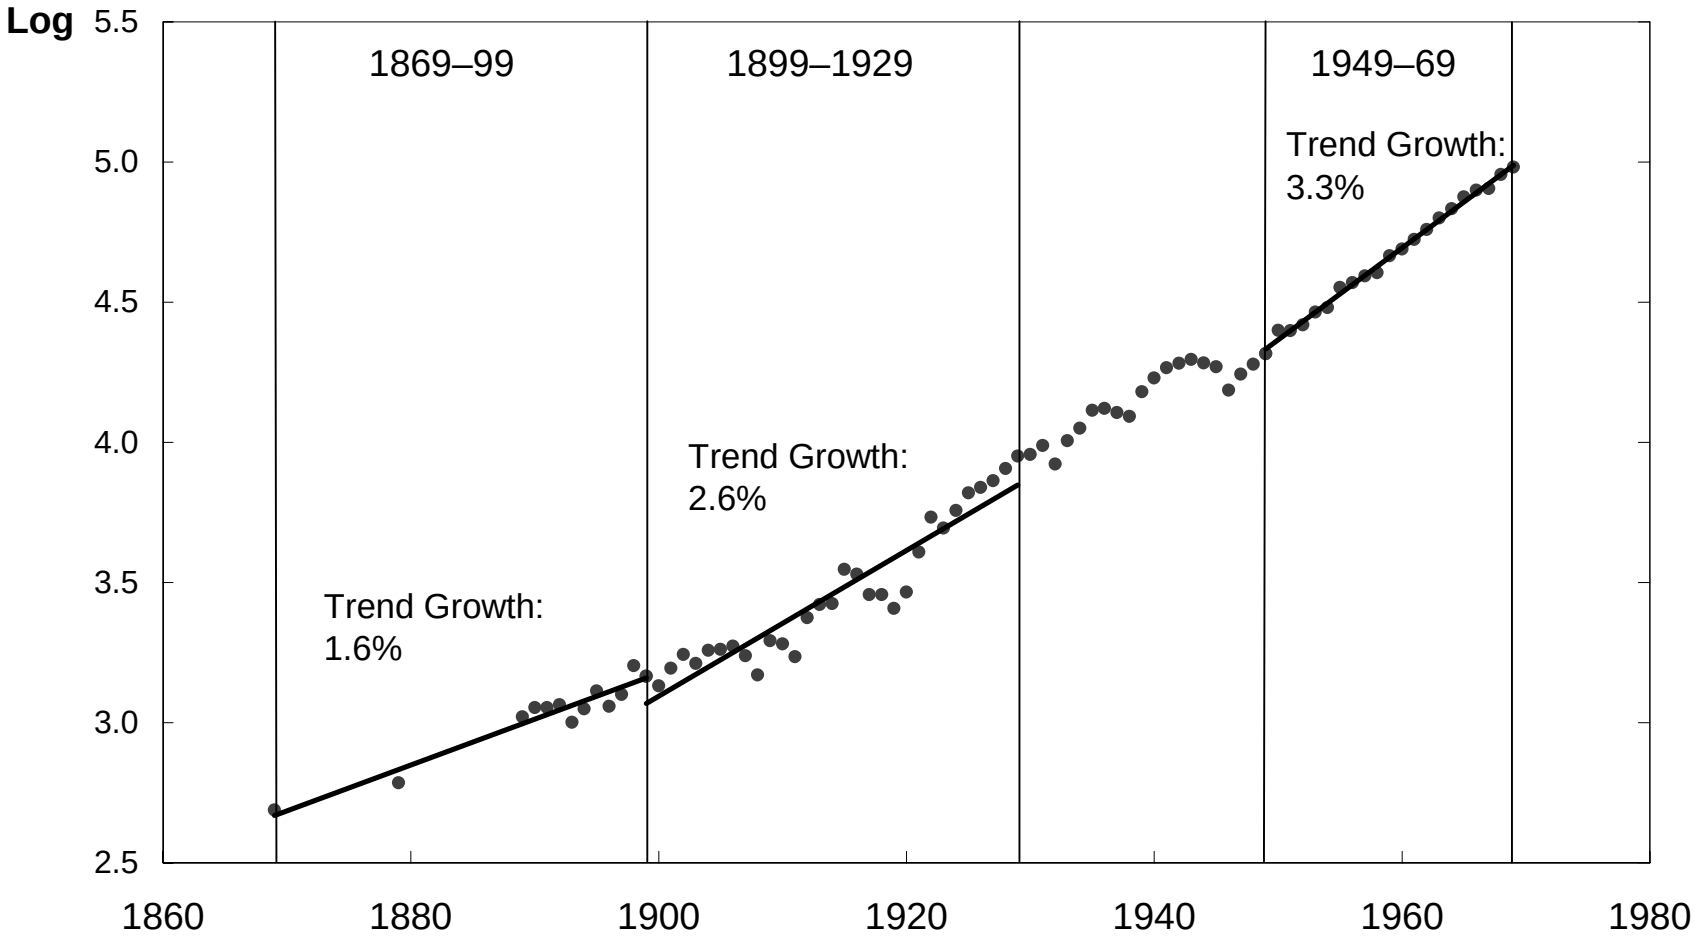

Source: U.S. Department of Commerce (1973)

Figure 3 **The Life Cycle of Plants in the Model**

Productivity of Two Plants vs. That of Frontier Blueprint Over Time

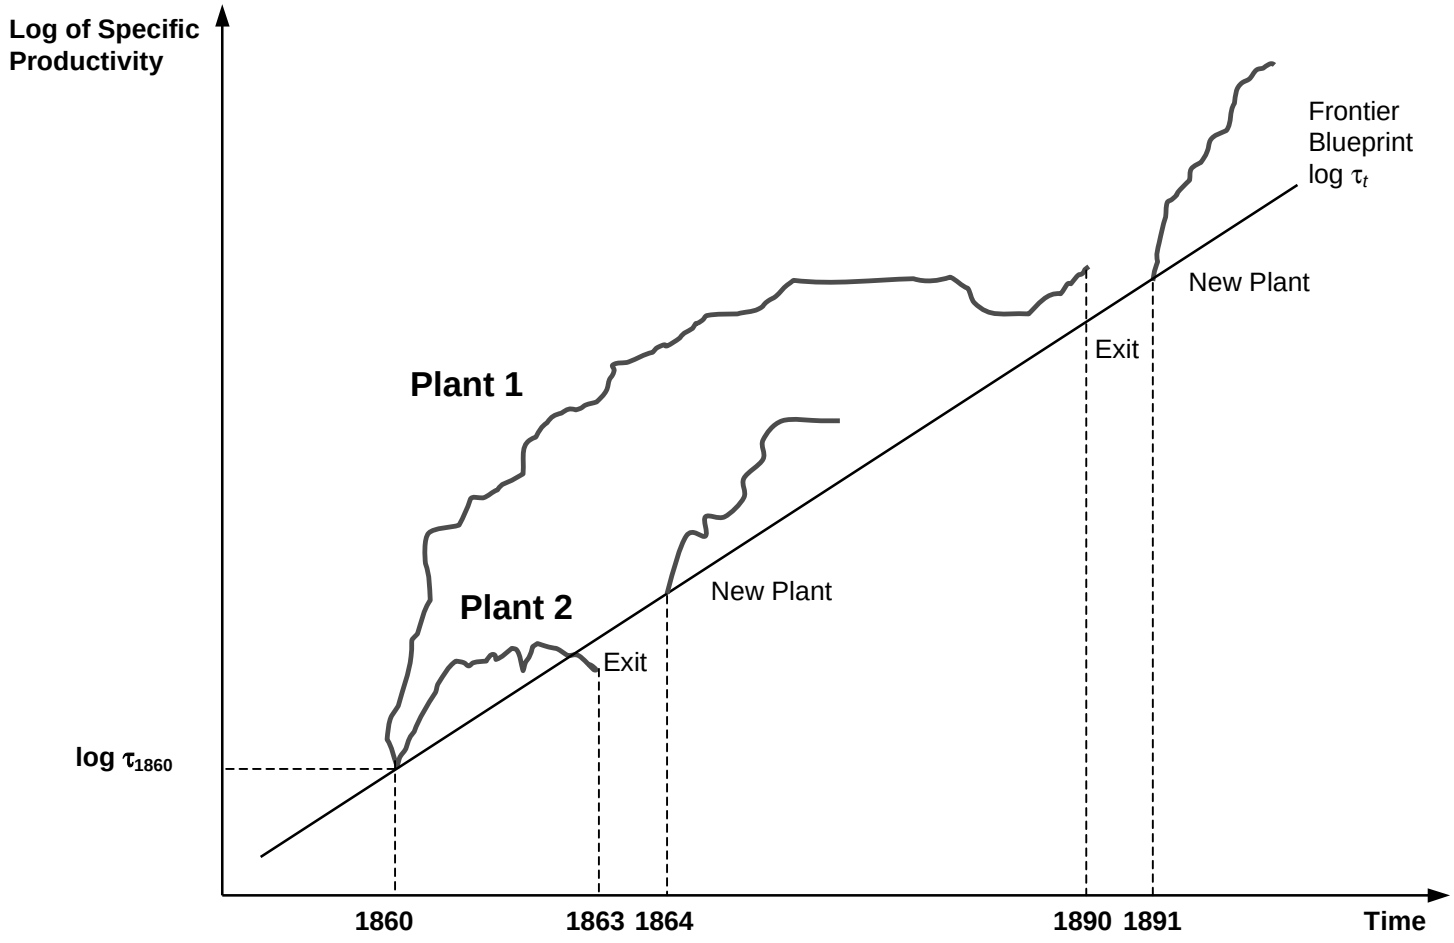

Figures 4–5 **Evidence That Plants Grow Relatively as They Age**

Figure 4 **Average Employment Share of One-Year Plant Cohorts**

Percentage of Total Workers Employed by U.S. Manufacturing Plants
of Various Ages in 1988

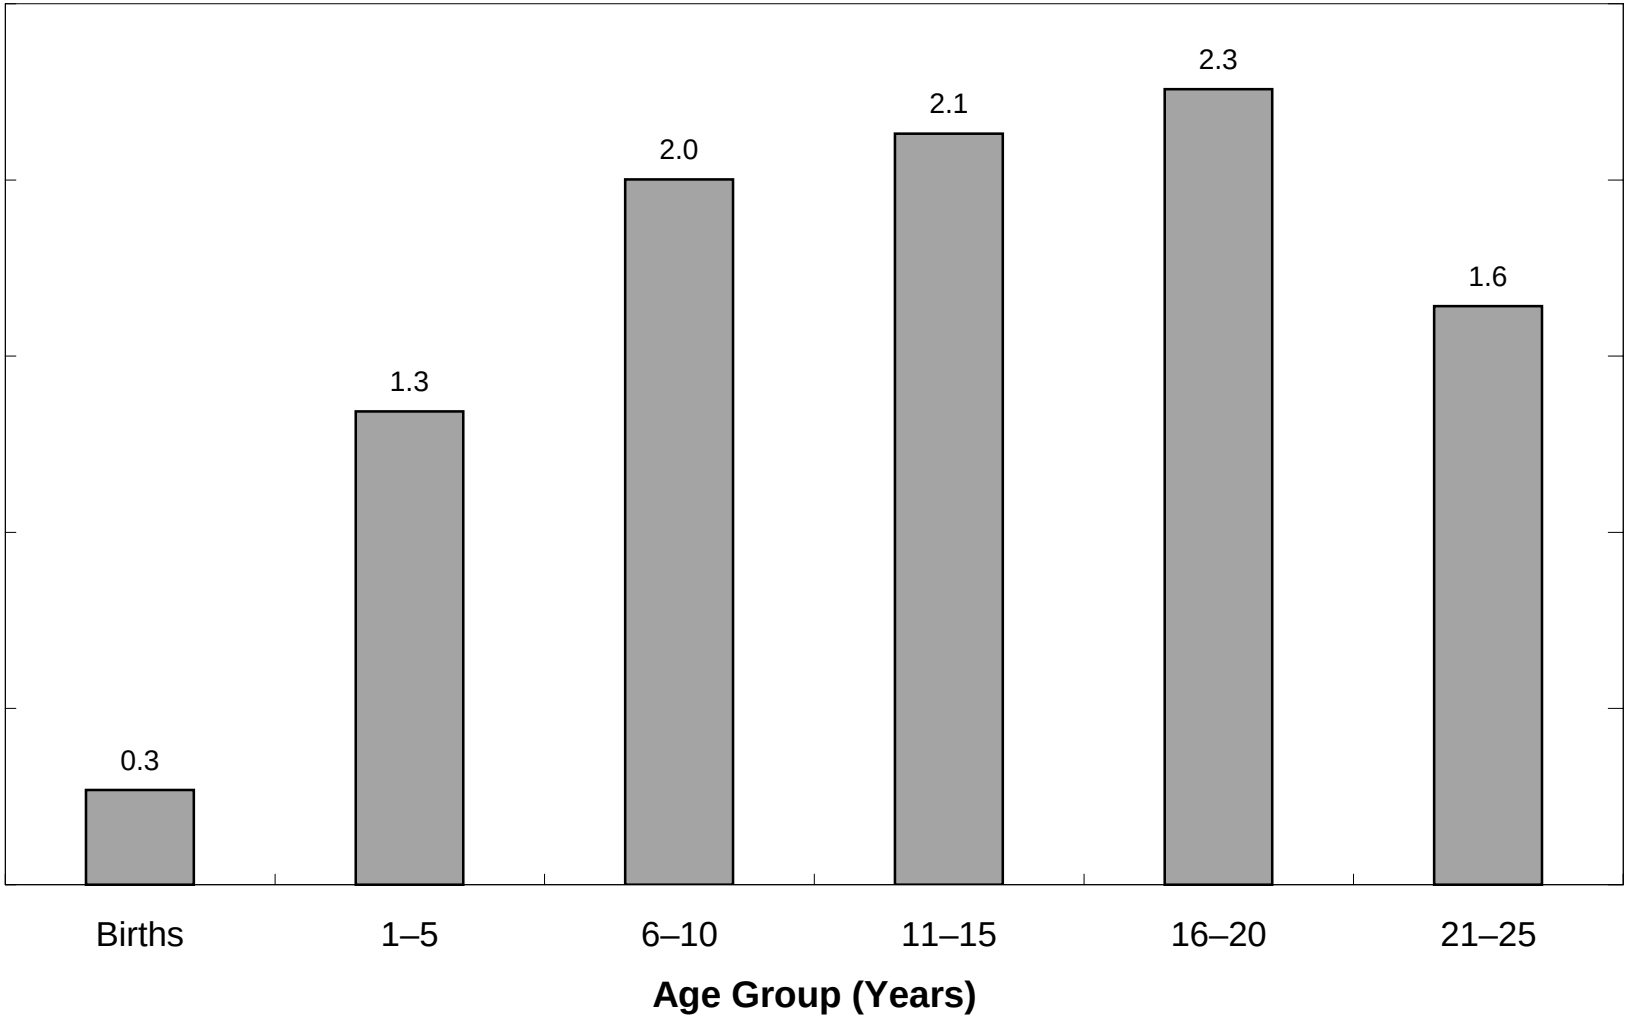

Source: U.S. Census Bureau's Longitudinal Research Database

Figure 5 **Average Hours Worked in Five Plant Cohorts**

Ratio of Average Hours in Each Plant Cohort to Average Hours in Total U.S. Manufacturing, 1967–87

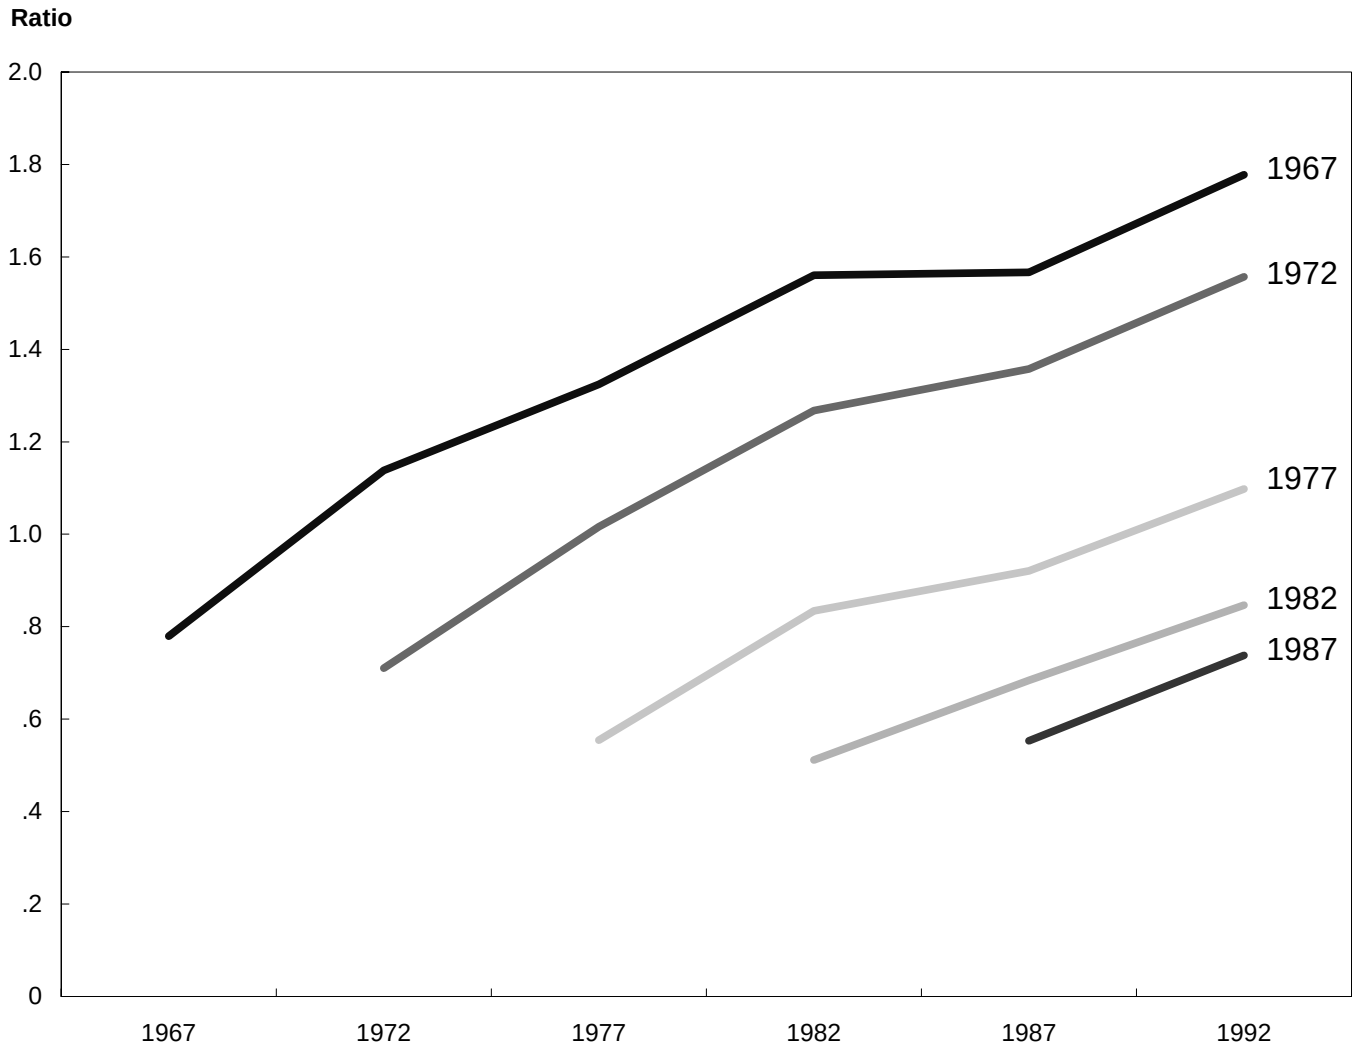

Source: Jensen, McGuckin, and Stiroh (2001)

Figure 6 **The Constant Relationship Between a Plant's Productivity and Its Age and Size**

Among U.S. Manufacturing Plants, 1972–86

A. Productivity By Age (Number of Years Operating)

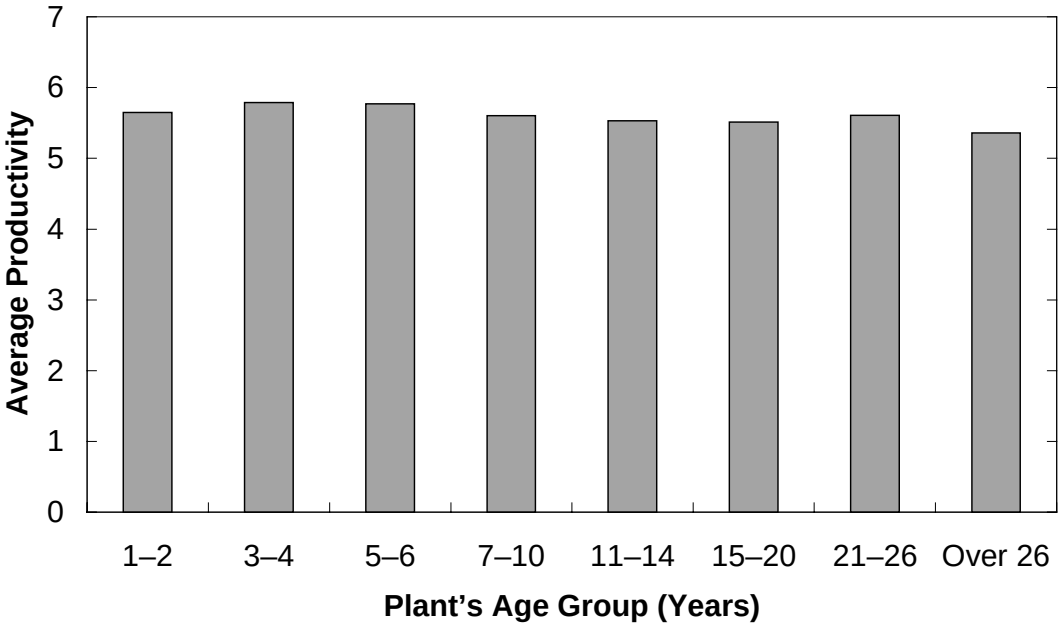

B. Productivity By Size (Number of Workers Employed)

Source: Bartelsman and Dhrymes (1998)

Figure 7 **The Model's Reproduction of U.S. Employment Data**

Each Statistic as a Percentage of Total U.S. 1988 Employment
for Manufacturing Plants of Various Ages

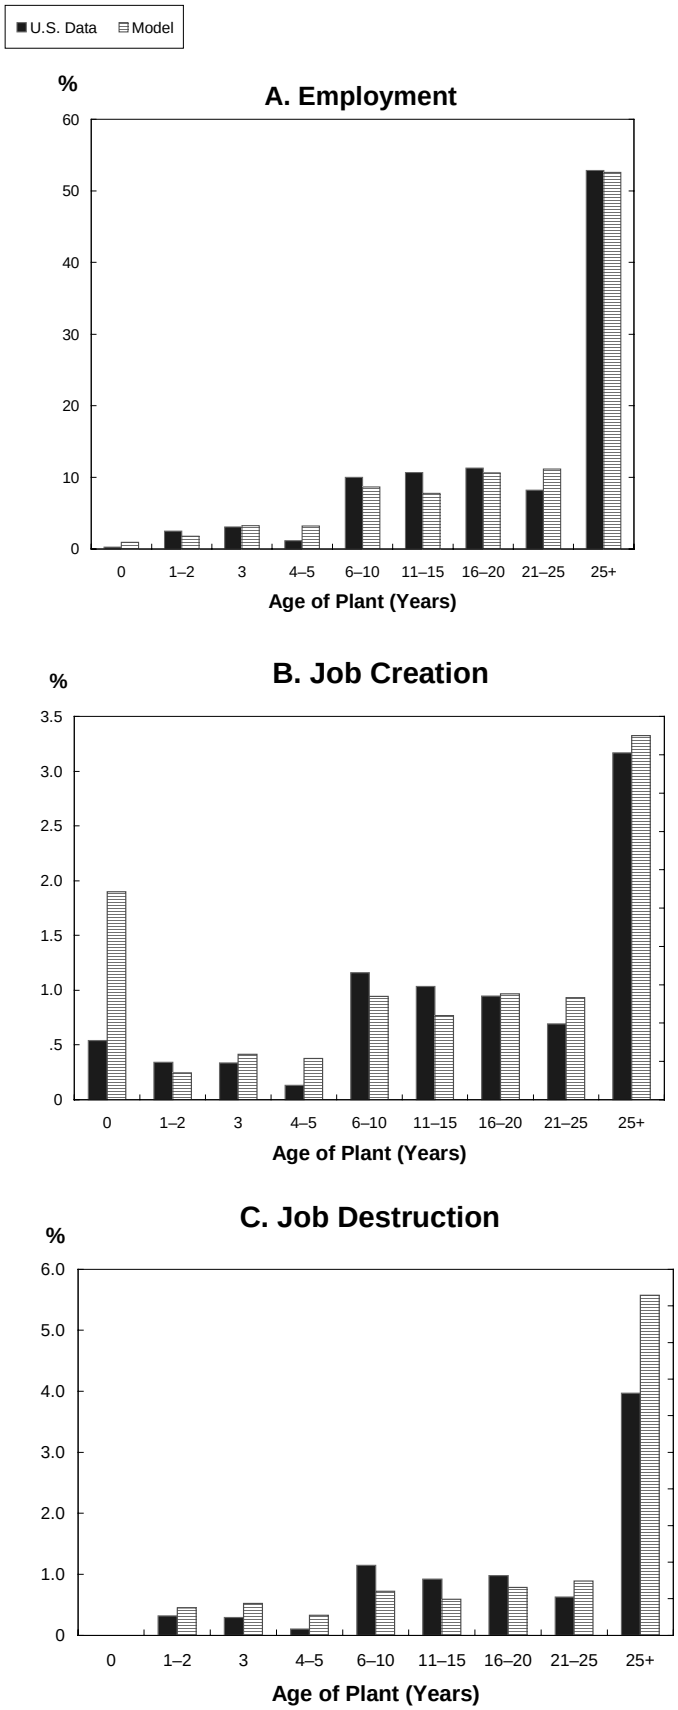

Source of U.S. data: U.S. Census Bureau's Longitudinal Research Database

Figure 8 **Characteristics of Model Micro Parameters on Learning**

Mean and Standard Deviation of Shocks to Plant Size and Productivity
by Age of Plant

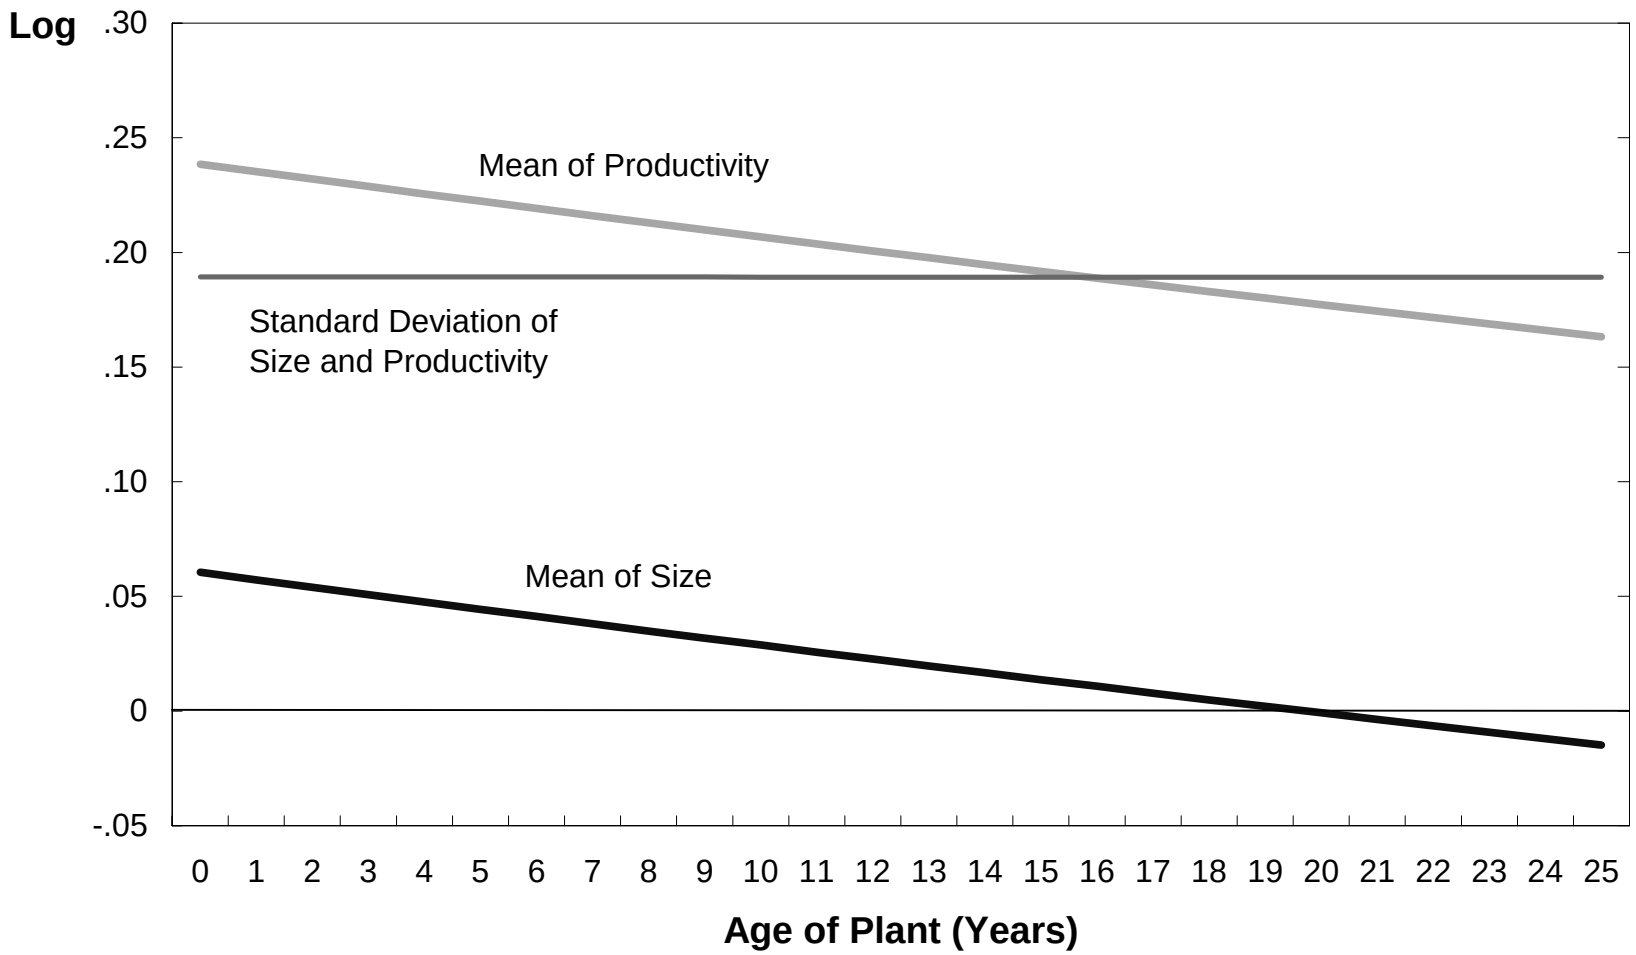

Figures 9–10 **The Model's Reproduction of the Gradual Transition to a New U.S. Economy**

Figure 9 **A Slow Diffusion of New Technology . . .**

Model's Predicted % of Output Produced in Plants with New Blueprints
vs. Actual % of Horsepower Provided by Electric Motors in U.S. Manufacturing Plants, 1869–1939

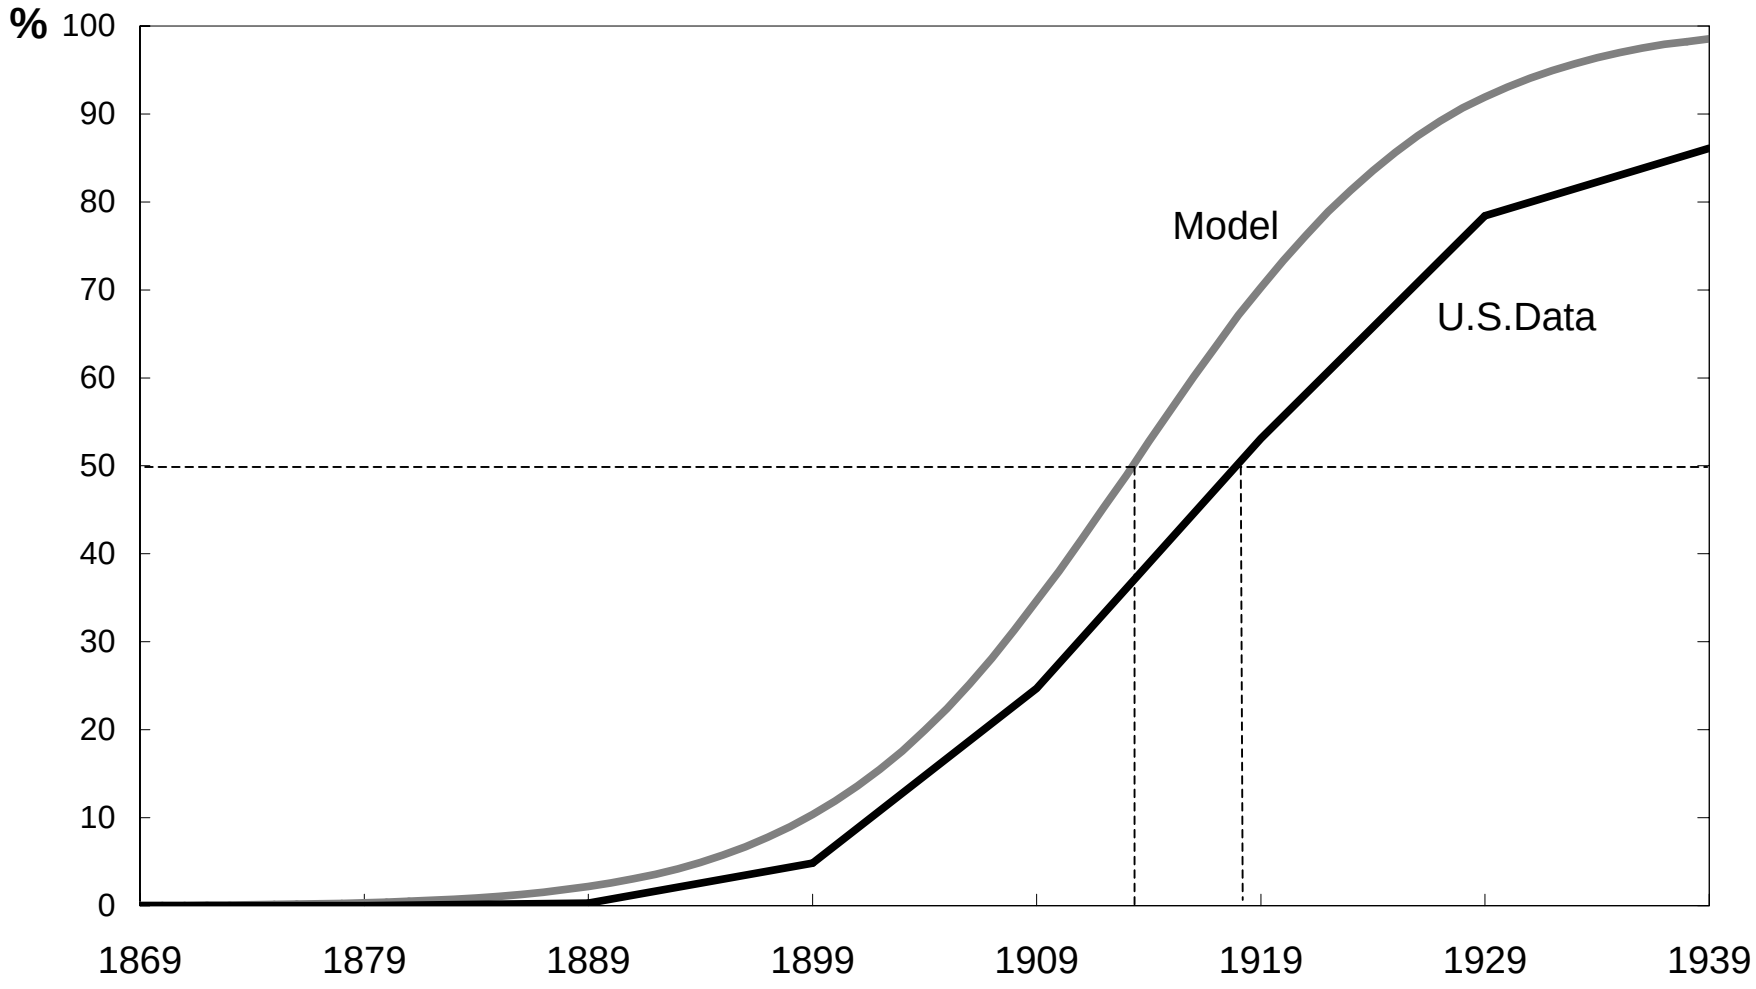

Source of U.S. data: Devine (1983, p. 351, Table 3)

Figure 10 **. . . And a Slow Increase in Productivity Growth**

Log of Output per Hour and Trend Growth in Manufacturing in the Model and the U.S. Data, 1869–1969

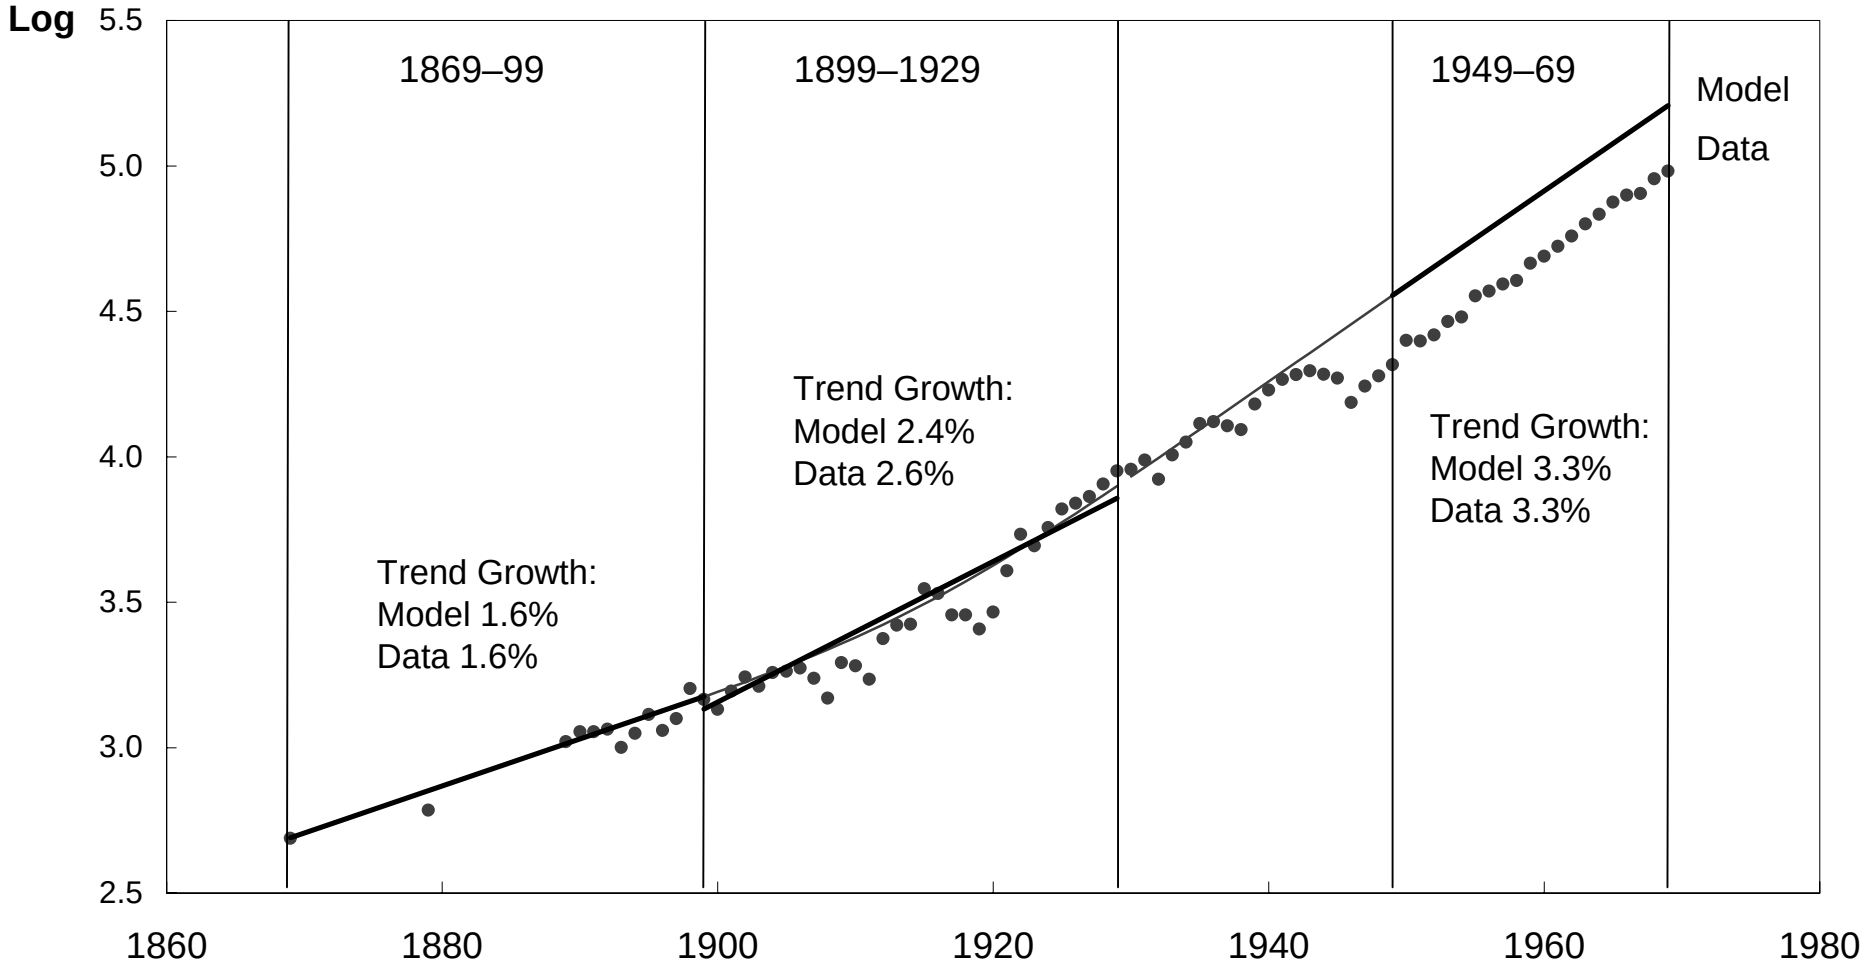

Source of U.S. data: U.S. Department of Commerce (1973)

Figure 11 **Initial Growth Rates Matter for the Speed of Transition**

Model's Predictions, With Various Initial Steady-State Growth Rates,
for the Time Until Diffusion Reaches 50%, the Time Until Productivity Growth Increases 1%,
and the Size of the Stock of Built-Up Knowledge in the Old Steady State as a % of That in the New

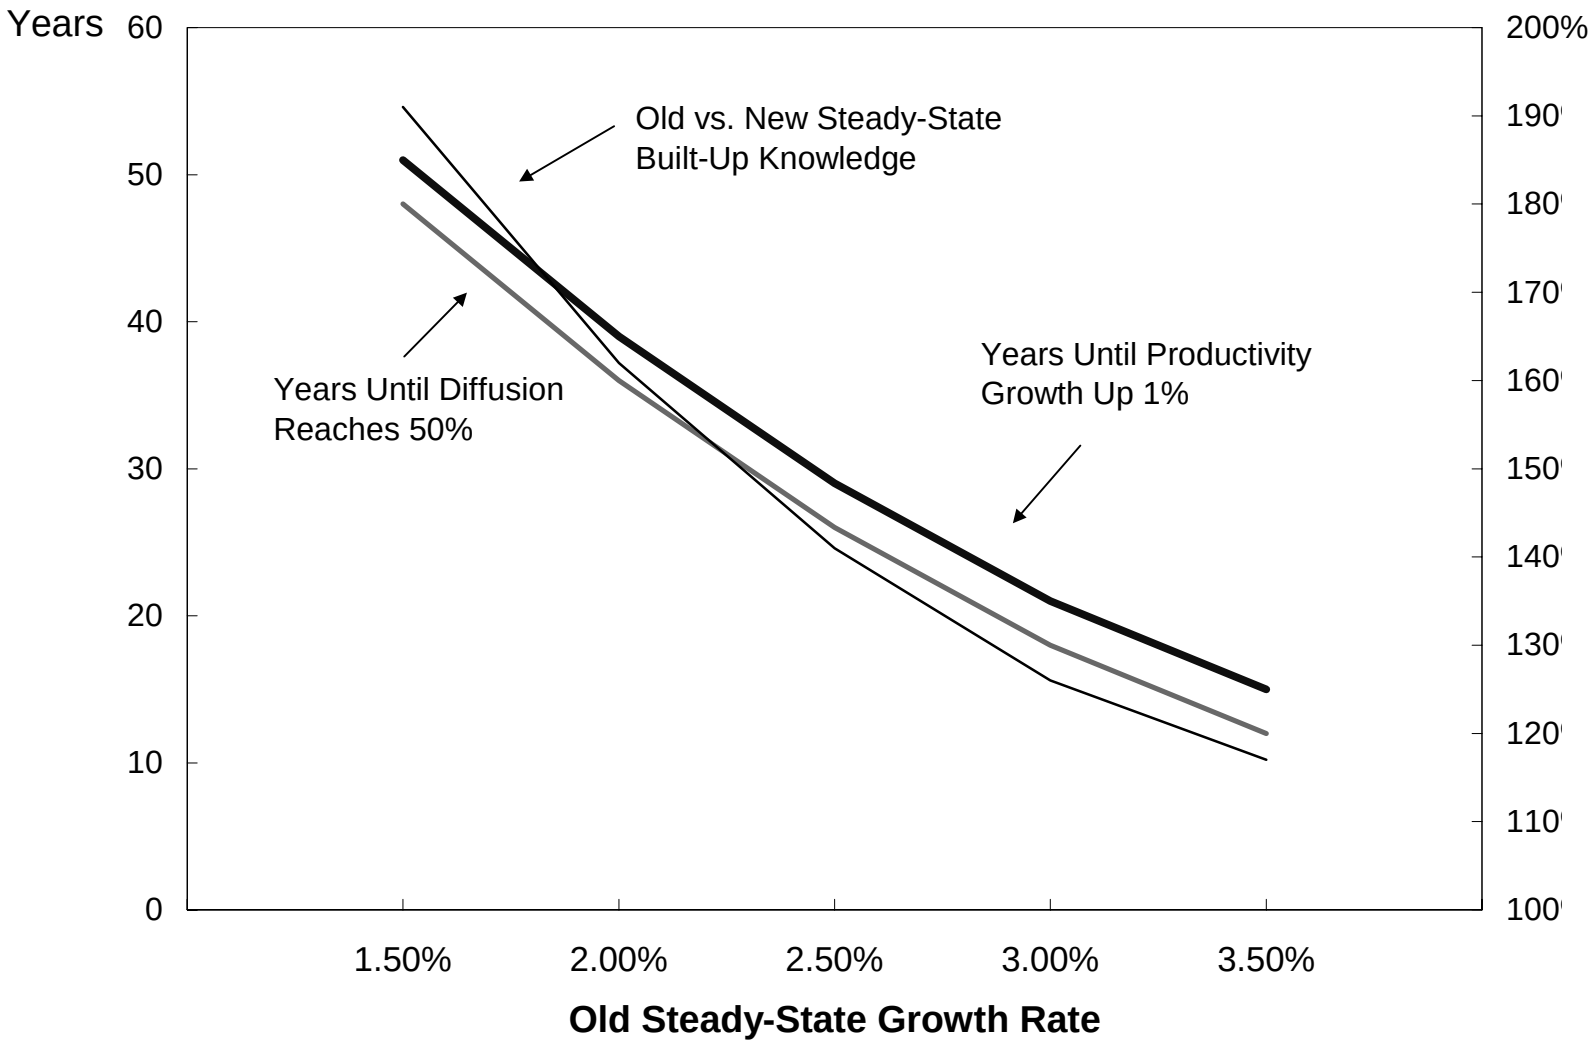

Figure 12 **Evolution of Power Distribution in Manufacturing Establishments**

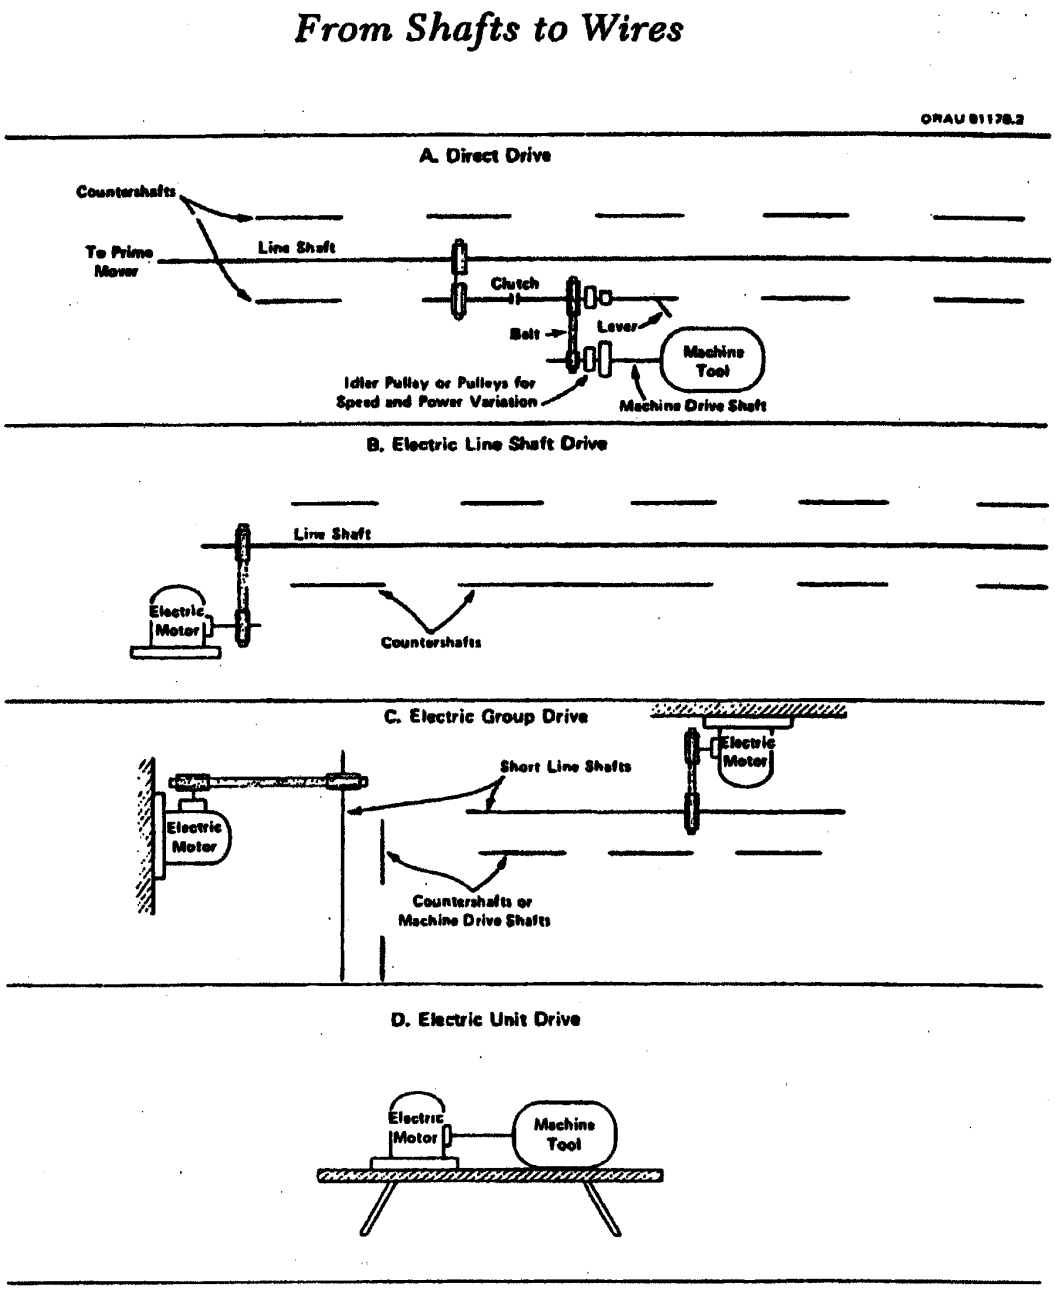

Source: Devine (1983, p. 353, Figure 2)

Figure 13 **Chronology of Electrification of Mechanical Power in Industry**

A. Methods of Driving Machinery
B. Key Technical and Entrepreneurial Developments

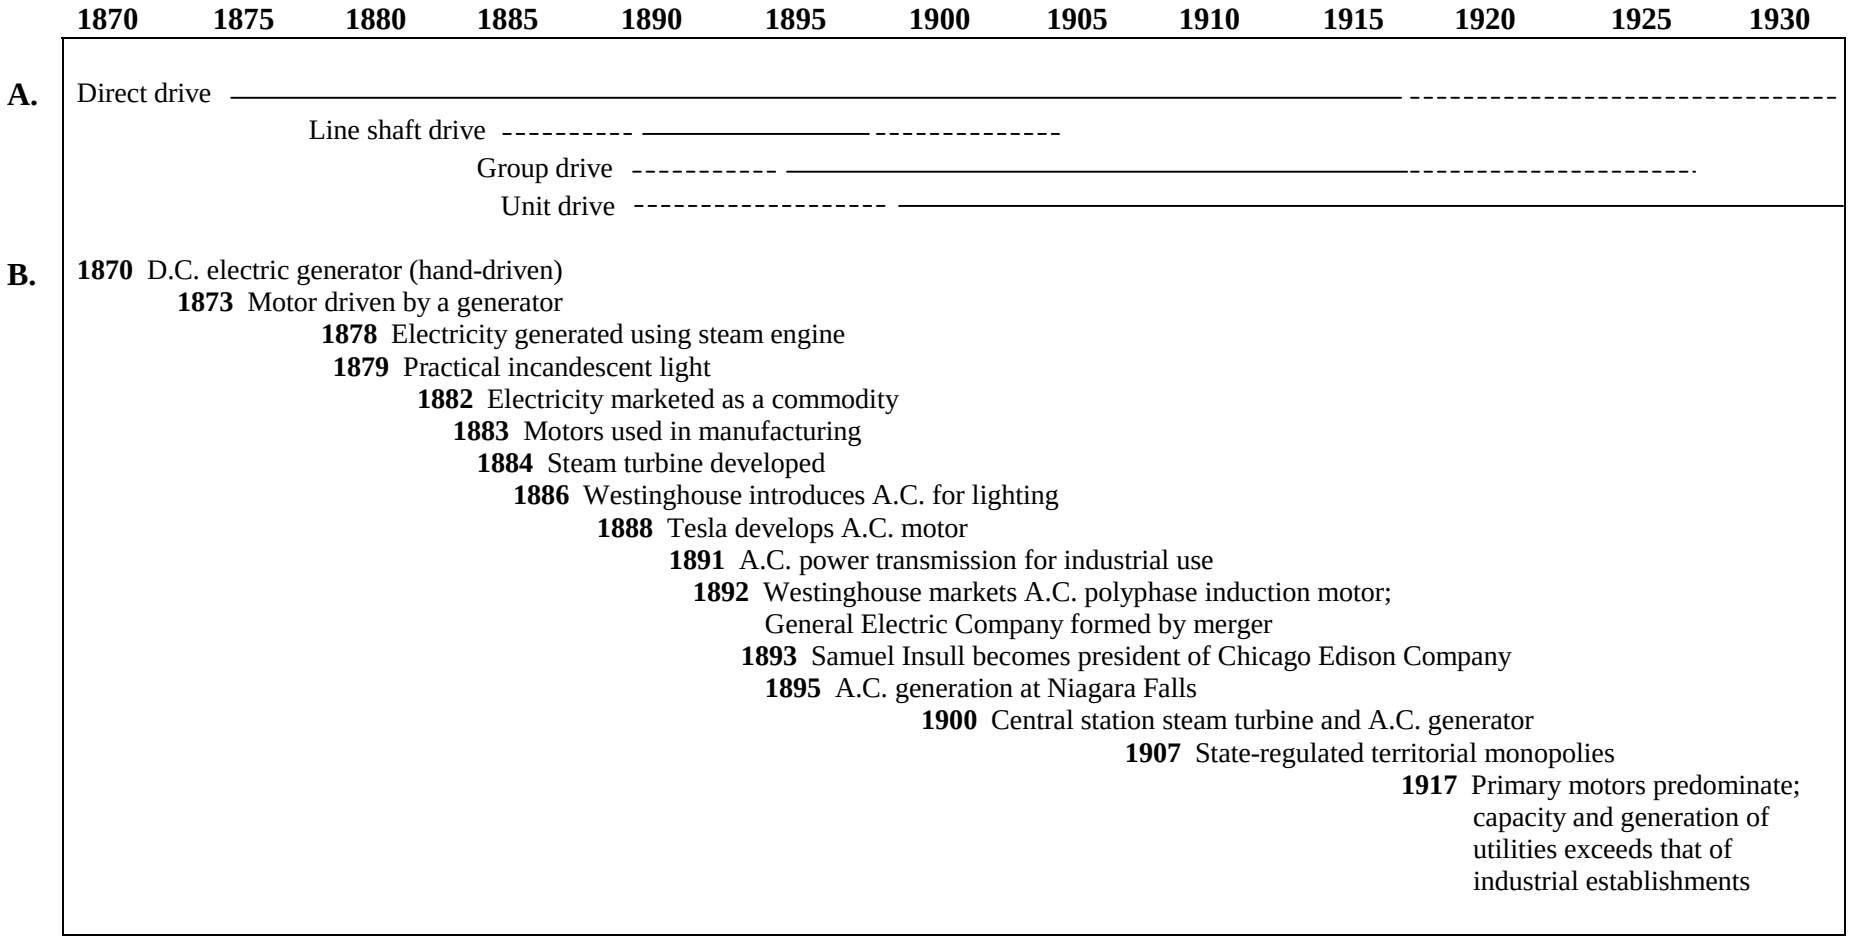

Source: Devine (1983, p. 354, Figure 3)

Figure 14 **Milestones in the Evolution of Production Organization**

1900	1910	1920	1930	1940	1950	1960	1970	1980	1990

1900 Machines arranged along line shafts

1910 Machines grouped by operation or product

1913 Flow-line assembly (Ford Motor Company)

1920 Assembly line common

1925 Large metal-working transfer machine

1930 Transfer machine for engine manufacture

1941 80-station transfer machine. Inflexible

1948 Word “automation” first used. Transfer machine common

1957 Greater mechanical integration (link lines and centralized control stations). Limited flexibility

1969 Highly developed automatic transfer machines. Improved flexibility

1980s Computerized materials-handling devices; flexible manufacturing systems

1985 Modular assembly (automatic guided vehicles)

1988 Computer-integrated manufacturing

Source: Devine (1990)