Technological Revolutions

By FRANCESCO CASELLI *

In skill-biased (de-skilling) technological revolutions learning investments required by new machines are greater (smaller) than those required by preexisting machines. Skill-biased (de-skilling) revolutions trigger reallocations of capital from slow- (fast-) to fast- (slow-) learning workers, thereby reducing the relative and absolute wages of the former. The model of skill-biased (de-skilling) revolutions provides insight into developments since the mid-1970's (in the 1910's). The empirical work documents a large increase in the interindustry dispersion of capital-labor ratios since 1975. Changes in industry capital intensity are related to the skill composition of the labor force. (*JEL* E23 J31 O33)

Income and wage inequality is rising in the United States, as well as in several other countries. Started in the 1970's, this large increase in earnings dispersion results from both absolute gains at the top, and absolute losses at the bottom, of the wage distribution. In fact, even median real earnings have declined. Various empirical explanations have been proposed for these trends, and the jury is still out on the quantitative importance of the different hypotheses. However, a growing consensus attributes a significant role to skill-biased technical change. The Information-Technology Revolution is the obvious suspect.[1] Despite the intensity of the empirical debate, theoretical work on skill-biased technological change has—with few exceptions—lagged behind. In this paper I present a simple model of technological revolutions. I explore the interaction between skills and technology, and offer an interpretation for the changing nature of this interaction. In doing so, I propose an explanation for the recent changes in the wage structure, including the absolute decline in the wage of unskilled workers. Further, the model sheds light on the possible future behavior of the wage structure: will the trend towards greater inequality continue or will it reverse itself? Finally, I generate predictions for other macroeconomic developments that can be expected to accompany the current changes in the wage structure.

Discussions of skill-biased technological change usually focus on issues of substitutability between skilled and unskilled labor. Instead, I propose to focus on substitutability among technologies. A technology is a combination of machines of a certain type and workers who have the skills necessary to use them. A technological revolution is the introduction of a new type of machines. Machines of the new type are more productive than machines of preexisting types, but they can only be operated by workers who have developed a set of *machine-specific* skills. The acquisition of such skills is costly, and the labor force is heterogeneous in the cost of acquisition. The revolution is *skill biased* if the new skills are more costly to acquire than the skills required by preexisting types of equipment. The revolution is *de-skilling* if the new skills can be acquired at a lower cost than the skills associated with preexisting technologies. I will argue in this paper that possible examples of skill-biased technological revolutions are the steam engine, the dynamo, and—of course—information technology. One possible exam-

* Graduate School of Business, University of Chicago, Chicago, IL 60637. I thank Alberto Alesina, Gadi Barlevy, Robert Barro, Eli Berman, Steve Davis, Sven Feldmann, Oded Galor, Zvi Griliches, Larry Katz, Peter Klenow, David Laibson, John Leahy, Greg Mankiw, Jonathan Morduch, Steve Tadelis, Mehmet Yorukoglu, Alwyn Young, Jaume Ventura, two referees, and the many others who have provided comments and/or data.

[1] Alternative hypotheses include the globalization of the world economy, changes in labor-market institutions, and the appearance of winner-take-all markets. For recent surveys of this literature see the symposium in the Spring 1997 issue of the *Journal of Economic Perspectives.*

ple of a de-skilling technological revolution is the assembly line.

When the revolution is skill biased, workers with low learning costs start using the new, more productive machines. High learning-cost individuals remain attached to tools of older types. In equilibrium, rates of return on new and old types of capital must be equalized. I assume that each technology features diminishing marginal returns to capital. Therefore, arbitrage opportunities can only be removed by increasing the capital endowment of workers using the high-productivity new tools relative to workers using the low-productivity old tools. Hence, capital flows away from workers in low-skill technologies and towards those in the high-skill technology. With a lower capital-labor ratio, high learning-cost workers experience an absolute decline in their wages. Because they work with more productive machines, *and* because they have a higher capital-labor ratio, low learning-cost workers see their wages rise.

Whether this initial increase in inequality is short- or long-lived depends on the relative strength of a number of conflicting forces during the *diffusion* process. I show that on the postrevolution growth path wage differentials keep rising. This increases the fraction of the labor force for which it is profitable to learn the new skills. It is possible that this steady growth in the proportion using the new tools will lead to universal adoption of the leading-edge technology. When this happens wage inequality declines. However, it is also possible for the economy to reach a steady state in which the highest learning-cost individuals have not upgraded their skills. In this case, the labor market remains indefinitely split into two tiers, wages are unequal, and there is incomplete adoption of the new production methods even in the long run.

When the revolution is de-skilling, a large number of workers can profitably embrace the new technology. As in the skill-biased case, the revolution increases the gains from using the most advanced technology. The difference is that the number of users of the high technology increases, while in skill-biased cases it shrinks. I show that this fact can lead to a fall in inequality during de-skilling revolutions. Furthermore, de-skilling revolutions always lead to the complete and rapid abandonment of at least some of the older technologies. This is in contrast with skill-biased revolutions, in which the diffusion process is generally characterized by a protracted phase of continued investment in the old technology, and sometimes by incomplete upgrading even in the long run.

When I survey the evidence I find that my description of a skill-biased technological revolution is consistent with many of the known facts on the recent changes in the wage structure. In addition, I also document some previously unknown facts about the interindustry dispersion of capital per worker. Specifically, I show that—after the mid-1970's—there has been a dramatic increase in the inequality of average capital-labor ratios across industries. In addition, industries experiencing relatively large gains in the capital-labor ratio also experienced relatively large gains in average wages, as well as relatively large increases in the proportion of skilled (nonproduction) workers in total employment. Finally, capital-labor ratios increased the most in industries that had high *initial* (1975) average wages and nonproduction worker shares. These facts are interesting in their own right, but I also argue that my theory constitutes one possible explanation for them.

This paper is concerned with those instances of technological progress that consist of the first-time appearance of a new family of tools, whether learning how to use them is harder (skill-biased revolution) or easier (de-skilling revolution). This paper is not concerned, instead, with those instances of technological progress that consist in improved versions of already existing tools: say, a more durable hammer. In these cases, where progress is *incremental,* the set of actions required of the worker to employ the improved tool is exactly the same as the one required to use the previous version, so that no acquisition of new skills is necessary. While revolutionary technological change modifies the wage structure, incremental technological change does not. My working hypothesis is that technological progress has been predominantly incremental in the 1950's and 1960's,[2] and predominantly

[2] In partial support of this conjecture, one may note that the ''hot'' industry in the 1950's and 1960's was the

revolutionary (of the skill-biased variety) in the 1970's and 1980's. The revolution in question is, of course, the Information-Technology Revolution.[3]

Two other papers share a similar premise. In Oded Galor and Daniel Tsiddon (1997), high- (low-) ability individuals have an incentive to work in the new (old) technology. At the same time, workers whose parents worked in the new (old) sector have an incentive to remain in the new (old) sector. The link between technological change and wage inequality depends on the life cycle of technology. In the early stages of a technological revolution the ability effect dominates, and this leads to increased income mobility and wage inequality. In the later stages the parental human-capital effect takes over, and income mobility and wage inequality decline. Jeremy Greenwood and Mehmet Yorukoglu (1997) emphasize the role of skilled workers in plantwide learning by doing. A technological revolution brings about both an acceleration in the rate at which the quality of new vintages of capital exceeds the quality of older vintages, and a decline in the initial level of "know-how" of newly born plants. Hence, when technological progress accelerates, new plants have to start from a lower level of productivity. Since the role of skilled workers is to improve technological know-how, their relative wages are bid up in the early stages of a technological revolution. It should be clear that these papers and mine study different channels by which technological change affects the wage structure, and should be considered complementary.[4]

This paper is also related to the literature on capital-skill complementarity (e.g., Zvi Griliches, 1969; Per Krusell et al., 1997), as it predicts that skilled workers will be endowed with more capital. Indeed, since the association between skill and capital is derived endogenously, my model may be interpreted as providing an explanation for capital-skill complementarity. Finally, the paper is related to the literature on the diffusion of new technologies (surveyed in Boyan Jovanovic, 1996), as it tracks how the share of output produced with the new type of machines increases over time. As mentioned, if learning costs are sufficiently high, this diffusion of the new technology can be slow and, in certain cases, it can grind to a halt before the older technology has been completely abandoned.

Section I presents the basic results for a simple case in which the economy experiences only one skill-biased technological revolution in its entire history. Section II generalizes to the case in which there are multiple revolutions. Section III shows how to extend the model to incorporate the case of de-skilling technological revolutions. Section IV reviews existing evidence and introduces new stylized facts on the interindustry distribution of capital per worker. Section V briefly explores historical antecedents. Section VI summarizes.

chemical industry. The so-called "chemicals revolution" allowed firms to replace a variety of inputs and products of organic origin with much cheaper inorganic materials (e.g., Edwin Mansfield, 1977; David C. Mowery and Nathan Rosenberg, 1989). While this replacement must have been accompanied by some revision in manufacturing practices, one may conjecture that the basic skills involved in handling the new materials were similar to those required to handle the old ones.

[3] In its *Surveys of Manufacturing Technology* (1988, 1993) the U.S. Bureau of the Census identified 17 new types of equipment, the appearance of which can be broadly thought of as giving rise to the Information-Technology Revolution. These are: computer-aided design (CAD); CAD-controlled machines; digital CAD; flexible manufacturing systems/cell; numerically controlled machines/computer-controlled machines; material working lasers; pick/place robots; other robots; automatic storage/retrieval systems; automated guided vehicle systems; technical data networks; factory networks; intercompany computer networks; programmable controllers; computers used on factory floors; automated sensors used on inputs; and automated sensors used on final products.

[4] Giovanni L. Violante (1997) also offers a complementary explanation for the increase in inequality *within* educational groups. In mine and the other contributions cited the arrival of new technologies is treated as exogenous. Daron Acemoglu (1998) provides a model in which the skill (or unskill) bias in technological change is endogenous, based on the relative supply of skilled and unskilled workers. Similarly, Huw Lloyd-Ellis (1999) formulates a model in which the rate of technical progress depends on the average quality of the workforce.

I. Technological Revolutions

A. *Growth-Theoretic Representation*

Revolutionary technological change is driven by substitution between *types* of capital. Types of capital differ in the set of skills required of the workers who use them. In any period output, Y_t, can be obtained using a number of different technologies. For example, production of, say, cars takes place both in plants where blue-collar workers use hand tools and in plants where engineers use robots. Technology j combines an amount K_t^j of capital of type j with an amount L_t^j of labor of type j, to produce $F^j(K_t^j, L_t^j)$ units of the final good. Labor of type j is labor endowed with the skills required to operate j-type machines. A technological revolution is the arrival of a new technology, i.e., of a new capital-skill *pair* capable of generating output. The aggregate production function is:

$$Y_t = \sum_{i=0}^{\infty} d_t^i F^i(K_t^i, L_t^i), \tag{1}$$

where

$$d_t^i = \begin{cases} 1 \ \textit{if} \ i \leq g(t) \\ 0 \ \textit{if} \ i > g(t) \end{cases} \tag{2}$$

and $g(t)$ is a nonnegative, nondecreasing function of time, to be specified later. Equations (1) and (2) imply that if $d_t^i = 1$ technology i is a feasible way of producing output in period t. The function $g(t)$ determines which technologies are feasible at a given date. Since $g(t)$ is nondecreasing, once a technology has come into existence it will always remain feasible (though it may not be used). (Unit) increments in $g(t)$ introduce new technologies. For example, if $g(t) = t$ there is a new technology being introduced in every period. Of course, $g(t)$ can follow a stochastic process. Period T is a period of revolutionary technological change if and only if:

$$\sum_{i=0}^{\infty} d_{T+1}^i = \sum_{i=0}^{\infty} d_T^i + 1.$$

The production functions are also indexed by i. I assume

$$F^i(K_t^i, L_t^i) = (A)^i (K_t^i)^{\alpha} (L_t^i)^{1-\alpha}, \tag{3}$$

where $0 < \alpha < 1$ and $A > 1$ are exogenous parameters. The specification in (3) captures the progressive nature of technological revolutions: it says that each newly introduced skill-type, capital-type pair is more productive than its predecessors, for given quantities of inputs.[5]

One unit of output can be converted into either one unit of the consumption good, or one unit of any of the existing types of capital. There are no available forms of investment other than capital goods. All markets are perfectly competitive. All capital and labor are owned by the household sector, which rents capital and labor of type i to firms at rates R_t^i and W_t^i, respectively. The consumption good is the numeraire. These assumptions lead to the usual factor-pricing equations:

$$R_t^i = \alpha A^i \left(\frac{L_t^i}{K_t^i}\right)^{1-\alpha} \tag{4}$$

$$W_t^i = (1-\alpha) A^i \left(\frac{K_t^i}{L_t^i}\right)^{\alpha}$$

for every i such that $d_t^i = 1$. I assume that all types of capital fully depreciate in one period. In Appendix C I discuss the consequences of slower depreciation.

Demographics are described by an overlapping-generations model. Each generation lives two periods, and is composed of a continuum of agents with aggregate measure 1. The objective of a member of the generation born at time t (generation t) is to maximize:

$$\log(C_t^y) + \beta \log(C_{t+1}^o),$$

where C_t^y (C_{t+1}^o) is her consumption when young (old), and $\beta < 1$ is an exogenous

[5] Incremental technological progress is easily incorporated into this representation by making A a function of time. If $g(t) = 0$ for every t we then have the standard neoclassical growth model.

parameter. The only endowment of each agent is one unit of labor, which is supplied inelastically when young. Before joining the labor market a worker must choose which type of machine-specific skills to acquire. Training into skills of type i comes at a consumption cost σ_i, and gives access to machines of type i. The workforce is heterogeneous with respect to cognitive ability and/or access to credit and education. In other words, individuals differ in learning costs σ_i. The latter are assumed to be distributed among the members of each cohort with cumulative distribution functions $\Phi^i(\sigma_i)$.[6]

B. *The Thought Experiment*

A simple special case provides the basic insights on skill-biased technological revolutions. Imagine a stationary economy in which, for a long time, only one kind of tools (hammers) has been available to be combined with labor in order to produce a final good. However, at some date T, a new and different brand of tools (robots) is invented to produce the same goods. The objective is to characterize the dynamics of the wage structure around date T. A useful simplification towards this goal is to pretend that time T is the only period of revolutionary change. In other words, I assume that $g(t)$ follows the process

$$(5) \qquad g(t) = \begin{cases} 0 & \text{for} \quad t \leq T \\ 1 & \text{for} \quad t > T \end{cases}.$$

This means that before the revolution only labor and capital of type 0 are used, and that the revolution adds machines and skills of type 1 to the menu of technologies. I also assume, without further loss of generality, that the skills required by technology 0 can be learned without cost (i.e., $\Phi^0(0) = 1$). Finally, I assume that $\Phi^1(\cdot)$ is uniform on the support $[0, \bar{\sigma})$. In Section II I return to the full model with multiple revolutions.

[6] An unpublished Appendix discusses the consequences of expressing σ_i as a time cost, instead of a consumption cost.

C. *Prerevolutionary Period*

Prior to date T all workers use the same type of tools. The stock of these tools at time t is denoted by K_t^0. Individual budget constraints are given by:

$$C_t^y = W_t^0 - K_{t+1}^0$$

$$C_{t+1}^o = R_{t+1}^0 K_{t+1}^0.$$

The solution to the consumer's problem gives rise to the savings function

$$K_{t+1}^0 = \frac{\beta}{1+\beta} W_t^0.$$

If there were no revolutions this economy would converge to the standard steady state of the overlapping-generations economy with logarithmic utility and Cobb-Douglas production function. Let an overline denote the value of a variable in the steady state of the one-type-of-capital economy. For T large enough, the economy will find itself in a neighborhood of this steady state when the revolution arrives. Hence, there is no loss of generality in assuming $K_t^0 = \bar{K}$ for $t \leq T$. The (well-known) formula for $\bar{K}$ is $(1 - \alpha)^{[1/(1-\alpha)]}[\beta/(1 + \beta)]^{[1/(1-\alpha)]}$. The formulas for $\bar{W}$ and $\bar{R}$ (hence, W_t^0 and R_t^0, $t \leq T$) follow immediately.

D. *Revolutionary Period*

At time T, a new generation of tools is invented, and will potentially become available for production in period $T + 1$. The sequence of events is as follows. At the end of period T the young decide how much to save, and how to allocate their savings between type-1 and type-0 capital. At the beginning of period $T + 1$ the young decide whether or not to acquire type-1 skills. Generation T's budget constraints are:

$$C_T^y = \bar{W} - K_{T+1}^0 - K_{T+1}^1$$

$$C_{T+1}^o = R_{T+1}^0 K_{T+1}^0 + R_{T+1}^1 K_{T+1}^1.$$

Notice that the new machines are not yet ready to be used in production at date T, so that labor income is still $\bar{W}$. From the consumer's problem the supply of capital for period $T+1$ follows the pattern:

$$(6)\qquad K^1_{T+1}\begin{cases} =\bar{K} & \text{if}\quad R^0_{T+1}<R^1_{T+1}\\ \in[0,\bar{K}] & \text{if}\quad R^0_{T+1}=R^1_{T+1}\\ =0 & \text{if}\quad R^0_{T+1}>R^1_{T+1}\end{cases}$$

and $K^0_{T+1}=\bar{K}-K^1_{T+1}$. In (6) I have exploited the fact that, with log utility, total savings will not depend on the interest rate. The top and the bottom lines correspond to the cases in which only one technology is used in period $T+1$ production, while the middle line describes the case in which both technologies are employed.

If any investment takes place in the new brand of machines, young workers in period $T+1$ have the option to gain access to these machines by paying the learning cost. From the assumptions on the learning-cost distribution and preferences, labor supplies in period $T+1$ are given by:

$$(7)\quad L^1_{T+1}=\Phi^1(W^1_{T+1}-W^0_{T+1})$$
$$=\max\left\{0,\min\left[1,\frac{(W^1_{T+1}-W^0_{T+1})}{\bar{\sigma}}\right]\right\}$$

and $L^0_{T+1}=1-L^1_{T+1}$. The first equality in (7) says that all workers for whom the learning cost is less than the wage differential will join the leading-edge sector, while those for whom the wage differential is not enough to pay for the learning cost will keep working with the old technology. The second equality follows from the assumption that the learning-cost distribution is uniform on $[0,\bar{\sigma})$.

In Appendix A I show that there is a unique equilibrium in period $T+1$. In particular, there is a cutoff level $\tilde{K}$ such that if $\bar{K}\geq\tilde{K}$ the equilibrium features immediate and complete abandonment of technology 0, while if $\bar{K}<\tilde{K}$ the equilibrium features partial adoption of the new technology, with both investment and labor split into 1 and 0 types. The cutoff $\tilde{K}$ is decreasing in A and increasing in $\bar{\sigma}$. Note that A measures the productivity gain associated with embracing the new technology, while $\bar{\sigma}$ indexes the cost of learning how to use it. Hence, for a given level of $\bar{K}$, if the "bang" A is large relative to the "buck" $\bar{\sigma}$ all workers start using technology 1 as soon as it becomes available. If, instead, the new technology generates a moderate productivity gain (relative to the learning cost), only a fraction of the workforce—those with the lowest learning costs—will upgrade their skills. It is the second, partial-adoption scenario that is of interest for our purposes, and I assume henceforth that parameter values are in the region giving rise to this type of equilibrium.

With that assumption, in the time $T+1$ equilibrium we have $L^1_{T+1}\in(0,1)$ and

$$(8)\quad W^0_{T+1}=[1+L^1_{T+1}(A^{[1/(1-\alpha)]}-1)]^{-\alpha}\times\bar{W}\quad<\bar{W}$$
$$W^1_{T+1}=[1+L^1_{T+1}(A^{[1/(1-\alpha)]}-1)]^{-\alpha}\times A^{[1/(1-\alpha)]}\bar{W}\quad>\bar{W}$$
$$R^0_{T+1}=R^1_{T+1}=[1+L^1_{T+1}(A^{[1/(1-\alpha)]}-1)]^{1-\alpha}\bar{R}\quad>\bar{R}.$$

In its first period, the technological revolution generates an increase in the skill premium. In fact, workers staying with the old machines not only lose in relative terms, but also in absolute terms: their wage falls relative to the prerevolutionary period. Skill-upgrading workers gain both in relative and in absolute terms. These developments are interesting because they mimic the available evidence on the recent evolution of relative and absolute earnings of skilled and unskilled workers.[7]

[7] Taken literally, the model implies that the skill premium increases from a value of 0. In Section II we will see that in the general case (with multiple revolutions) skill-biased technological revolutions further increase already-existing skill premia.

The intuition for these results is straightforward. If a fraction of the labor force remains unskilled, these workers can only be employed in the old sector in period $T + 1$. Hence, labor-market equilibrium requires that $K^0_{T+1} > 0$. This means that period-T young must be persuaded to invest in the old type of machines, as well as in the new. In turn, this implies that the rental rates must be equalized across the two technologies, yielding the third equation in (8). Now, if capital-labor ratios were the same, the marginal productivity of the new machines would be greater than the marginal productivity of the old machines by a factor of A. Hence, given decreasing marginal returns to capital, to equalize rates of return the capital-labor ratio in technology 1 must be greater than in technology 0. This inequality in capital-labor ratios also implies that the capital endowment of unskilled workers is actually lower than it was before the introduction of the new technology. For, the same amount of total capital, $\bar{K}$, previously distributed equally among all workers, must now be distributed unequally: more to the skilled and less to the unskilled. The only way to achieve this is for the unskilled to receive less than they used to. It is this fall in the unskilled capital-labor ratio that explains the fall in the unskilled wage.

Notice that workers upgrading their skills enjoy a twofold advantage. First, they work with inherently more productive machines, so that they would already be earning a premium if they had the same capital-labor ratio as the unskilled. We may term this the *direct effect.* Second, they have a higher capital-labor ratio, an *indirect effect* that works through the general equilibrium of the model. The literature so far has focused on the direct effect. But the direct effect alone cannot explain the absolute fall in wages at the bottom.

Suppose that we date the beginning of the Information-Technology Revolution in the mid-1970's. Since this is an overlapping-generations model, what I have termed the revolutionary period may roughly correspond to the last two decades. I conclude that the model can replicate the behavior of the wage structure in the recent past.[8] To see what may happen to the distribution of wages in the future, however, we need to study the model's dynamics in the postrevolutionary period.

E. *Postrevolutionary Period*

Define $K_t \equiv K^1_t + K^0_t$ as the total amount of capital (in units of the consumption good) in period $t \geq T + 1$. Now note that each period t there is born a new generation of workers who face the same skill-acquisition decision as did the young workers in period $T + 1$. Hence, the labor-supply function for period t is as in (7) (with t instead of $T + 1$). Similarly, the interindustry supply of capital, conditional on K_t, is still given by (6) (with K_t instead of $\bar{K}$). Therefore, conditional on K_t, all the results obtained above for the revolutionary period (which are conditional on the initial capital stock $\bar{K}$) extend to all subsequent periods. In particular, the same threshold level for K_t, $\tilde{K}$, determines whether production in period t features both technologies ($K_t < \tilde{K}$), or only technology 1 ($K_t \geq \tilde{K}$).

The relationship between the state variable K_t and the equilibrium allocation in period t is depicted in (K_{t+1}, L_{t+1}) space in (the left side of) Figure 1. The function $h(K_{t+1})$ describes how sector-1 employment varies as a function of the state variable K_{t+1}, in periods $t + 1$ in which the economy is using two technologies. In Appendix B I show that $h(0) = 0$, $\lim_{x\to\infty} h(x) = \infty$, $h'(\cdot) > 0$, and $h''(\cdot) < 0$. Hence, the larger the total capital stock K_{t+1}, the larger the fraction L_{t+1} of workers who embrace the advanced technology. However, L^1 is bounded above at 1. The figure identifies $\tilde{K}$ as the minimum level of the capital stock such that the whole labor force decides to become skilled.

[8] Since workers are born with no assets, and savings are proportional to wages, increased wage dispersion also implies increased income dispersion. The consumption of young unskilled workers also declines. However, since the interest rate increases, higher returns on savings may potentially compensate old unskilled workers for their lower level of accumulated savings. Hence, consumption of old unskilled workers, unlike wages, may in principle increase.

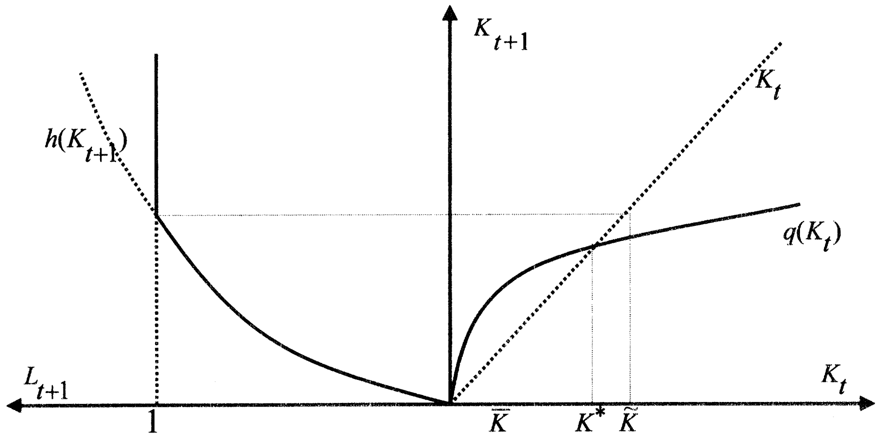

FIGURE 1. POSTREVOLUTIONARY DYNAMICS WITH $\tilde{K} > K^*$

We are left with the task of characterizing the dynamics of K_t after the revolution. Consider any period t in which the labor force is split into those operating the new brand of machines and those working with the old technology. One such period is, of course, period $T + 1$. A fraction L_t^1 has net labor income $W_t^1 - \sigma$, and hence will save amount $\beta/(1 + \beta)(W_t^1 - \sigma)$. The remaining $1 - L_t^1$ have labor income W_t^0 and save amount $\beta/(1 + \beta)W_t^0$. As a result, the total supply of savings on the capital market will be:

$$(9) \quad K_{t+1} = \frac{\beta}{1+\beta}\left[L_t^1 W_t^1 - \int_0^{W_t^1 - W_t^0} \sigma d\Phi^1(\sigma) + (1 - L_t^1)W_t^0\right].$$

Since K_t is the only state variable, and L_t^1, W_t^1, and W_t^0 are all functions of K_t, equation (9) implicitly defines a difference equation, $K_{t+1} = q(K_t)$, that governs the motion of K_t in periods in which the labor market is split between skilled and unskilled. In Appendix B I show that $q(0) = 0$, $\lim_{x\to\infty} q(x) = \infty$, $q'(\cdot) > 0$, and $q''(\cdot) < 0$.

The law of motion $q(K_t)$ is depicted in (K_t, K_{t+1}) space, together with a 45-degree line, in (the right side of) Figure 1. Notice that this path reaches a steady state when the capital stock is K^*. I assume that the initial capital stock, $\bar{K}$, is less than K^*.[9] Hence, Figure 1 depicts a situation in which the effect of a technological revolution is to start a growth process in total capital, from $\bar{K}$ towards K^*. As total capital increases, an increasingly large proportion of workers moves into sector 1. This process continues until the economy reaches the steady state K^*. In this steady state, the economy remains stuck (at least until the next revolution) with a two-tier labor market, two technologies, and wage inequality.

In Figure 1 the assumption is that $\tilde{K} > K^*$. This is not, however, the only possibility. Figure 2 also depicts the functions $h(K_{t+1})$ and $q(K_t)$, but it assumes that $\tilde{K} < K^*$. In this case the economy reaches $\tilde{K}$ in finite time. At this point all workers enter the skilled sector, i.e., $L^1 = 1$, and the skill premium disappears. After this has occurred, the economy is again a one-type-of-capital economy, and equation (9) ceases to describe its dynamics. The new law of motion for the capital stock is instead represented by the curve $\hat{q}(K_t)$, which displays further growth towards the steady state K^{**}. This steady state features a unique technology.[10]

[9] The case $\bar{K} > K^*$ can also easily be analyzed, but it is less interesting. It can also be ruled out for all reasonable values of α.

[10] For $K_t \geq \tilde{K}$ all the workers are using technology 1, so that $K_{t+1} = [\beta/(1 + \beta)][W_t^1 - \int_0^{\bar{\sigma}} \sigma d\Phi^1(\sigma)] \equiv$

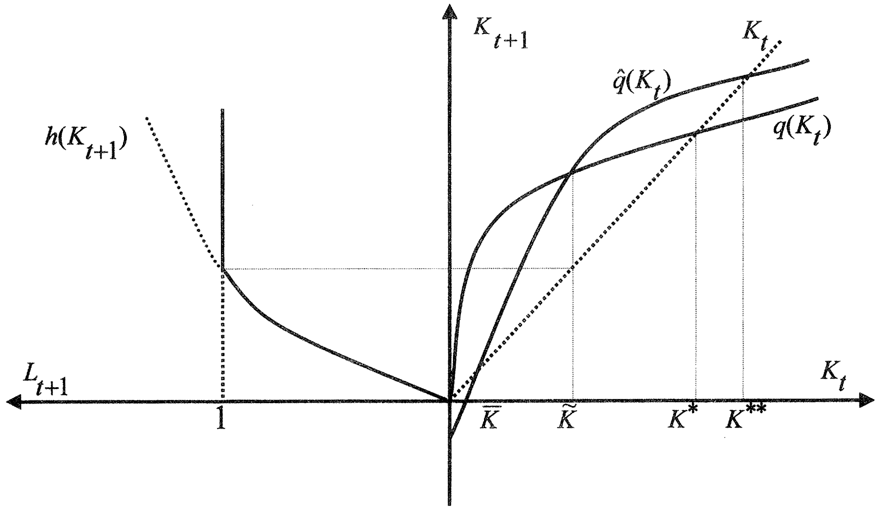

FIGURE 2. POSTREVOLUTIONARY DYNAMICS WITH $\tilde{K} < K^*$

In sum, if $\tilde{K} < K^*$ the economy fully adopts the new technology in finite time, while if $\tilde{K} > K^*$ the economy reaches the two-tier steady state and never fully upgrades. The intuition for these two possible outcomes is simple and, as before, it is driven by the behavior of the capital-labor ratios. The no-arbitrage condition must hold at all times, and it requires that the two capital-labor ratios be in constant proportions:

$$A^{[1/(\alpha-1)]}\frac{K_t^1}{L_t^1} = \frac{K_t^0}{L_t^0}.$$

Now, throughout the transition, total capital (the sum of the numerators above) increases while total labor (the sum of the denominators) is constant. Hence, in order to maintain the interest parity, both capital-labor ratios must increase. Increases in the two capital-labor ratios imply decreasing returns, giving rise to the concavities in the right sides of Figures 1 and 2. On the other hand, in periods in which both technologies are used, the labor-supply equation is simply [see (7)]:

$$(10) \qquad L_{t+1}^1 = \frac{(W_{t+1}^1 - W_{t+1}^0)}{\bar{\sigma}}.$$

The constant proportions of the capital-labor ratios imply that the wage *ratio* is also constant (the skill premium is always $A^{(1/1-\alpha)}$). This obviously means that, as rising capital-labor ratios raise both wages, the wage *differential* also increases. From the sector-1 labor-supply equation, each increase in the wage differential pulls further workers into the leading-edge sector. This gives rise to the increasing relationship depicted in the left side of the figures.

Notice that the model predicts continued investment in the old type of machines for many periods after the advent of the revolution. In fact, investment in old equipment

$\hat{q}(K_t)$. Using the factor-pricing equations and the assumptions on Ψ it is easy to see that $\hat{q}(K_t) = [\beta/(1+\beta)][(1-\alpha)AK_t^\alpha - 1/2\bar{\sigma}]$. It is straightforward to plot $\hat{q}(K_t)$ against K_t, as well as to show that $\hat{q}(\tilde{K}_t) = q(\tilde{K}_t)$. In addition, $\hat{q}(K_t) < q(K_t)$ for $K_t < \tilde{K}$ (Appendix B). Notice that, given an initial level $\bar{K}$, the dynamic equilibrium path of the economy is unique: for $K_t < \tilde{K}$ we have $K_{t+1} = q(K_t)$; for $K_t \geq \tilde{K}$ we have $K_{t+1} = \hat{q}(K_t)$. If $K^* < \tilde{K}$ there may be other two steady states to the right of K^* [imagine drawing $\hat{q}(K_t)$ in Figure 1]. The assumption $\bar{K} < \tilde{K}$ insures that the economy converges to K^*.

continues forever in the case depicted in Figure 1. Therefore, one can interpret the model as one of slow *diffusion* of new technologies: the presence of learning costs means that old and new technology can coexist for a long time.[11]

Further insights could be gained by studying alternative specifications for the learning-cost distribution. For example, these costs could be endogenous if there is a learning externality: if it is easier to pick up computer skills when there are many computer-literate individuals who can answer questions, and many machines to practice on, the cost of learning will shrink over time as the new technique diffuses throughout the economy. An unpublished Appendix—available on request—works out this case in detail.[12] Predictably, these learning externalities lead to faster adoption of the new technology, and increase the likelihood that the economy will fully upgrade in finite time.

More generally, the shape of the learning-cost distribution determines *the rate* at which workers move from sector 0 to sector 1 at different stages of the diffusion process. The same increase in $W^1 - W^0$ can generate widely different amounts of labor "migration" at different points of the cost distribution. In turn, such varying rates of migration can profoundly affect the rate of total capital accumulation. If, for example, the rate at which workers move from sector 0 to sector 1 increases very sharply over time, the curves in the right sides of Figures 1 and 2 may even become convex: rapidly increasing rates of reallocation of resources from the less to the more productive technology can more than compensate for the declining marginal products of capital generated by increasing capital-labor ratios.

[11] Other mechanisms for slow diffusion are explored in the technology adoption literature. One example is V. V. Chari and Hugo Hopenhayn (1991).

[12] The specific assumption is that the learning cost σ is uniformly distributed in the interval $[0, \bar{\sigma}/K_t^1]$, so that as the stock of capital of type 1 increases, and the new technology spreads, everybody's cost of learning endogenously declines.

II. Multiple Revolutions

By imposing condition (5) I have until now restricted exploration of the model to the special case in which there is only one revolutionary episode in the entire technological history of the economy. This section studies the general case, in which there are several revolutions.

Let $g(t) = i$ for $T_i < t \leq T_{i+1}$ where $\{T_j\}_{j=0}^{\infty}$ is an exogenously given sequence such that $T_j < T_{j+1}$. It will simplify the discussion, without substantial loss in generality, to think of technological revolutions as being separated by long intervals of time. Specifically, let us imagine that, as the jth technology becomes available for the first time, the transitional dynamics following the advent of technology $j - 1$ have already driven the economy into a neighborhood of the steady state to which the economy would converge should revolution j not occur (that such a steady state exists will be seen below). Then, we can treat each revolution as a shock that sets the economy in motion from one steady state to another.

Say, then, that the economy is in steady state at $t = T_j$ [hence, $g(t) = j - 1$], and assume that $j \geq 2$. Recall that this implies that the "menu" of technologies available to the economy at time T_j is composed of all capital-labor pairs i, such that $0 \leq i \leq j - 1$. We want to know, first, how the wage structure changes between time T_j and time $T_j + 1$. Second, we want to characterize the behavior of the wage structure during the transition following period $T_j + 1$. What makes the problem with $j \geq 2$ more complicated than the problem with $j = 1$ is that at time T_j the workforce can be spread across several technologies. Let $\tilde{L}_t \equiv \{L_t^0, \ldots, L_t^i, \ldots, L_t^{g(t)}\}$ be the distribution of employment among available technologies at time t. $\tilde{L}_{T_j}$ will generally be nondegenerate, although any number of its terms may be 0.

The initial impact on the wage structure of the arrival of technology j depends on a comparison of $\tilde{L}_{T_j}$ with the new distribution of employment, $\tilde{L}_{T_j+1}$. To see how, define the set of technologies being used at time t as $I_t \equiv \{i \mid L_t^i > 0\}$. At any time t, a no-arbitrage condition exists among any two types of equipment being used. This no-arbitrage condition can be stated as

$$A^{[(j-i)/(\alpha-1)]}\frac{K_t^j}{L_t^j} = \frac{K_t^i}{L_t^i} \qquad \forall i, j \in I_t.$$

Now define K_t, in a manner analogous to the previous section, as the total capital stock: $K_t \equiv \Sigma_{i \in I_t} K_t^i$. Then we can express wages in technology i as:[13]

$$(11)\quad W_t^i = \frac{A^{[i/(1-\alpha)]}(1-\alpha)K_t^{\alpha}}{[\sum_{i \in I_t} A^{[i/(1-\alpha)]}L_t^i]^{\alpha}} \qquad \forall i \in I_t.$$

(Note that the formulas for W_{T+1}^0 and W_{T+1}^1 in Section I constitute a special case of this.) Since we have assumed that the economy is in steady state at time T_j (implying $K_{T_j} = K_{T_j+1}$), the last expression shows that knowledge of $\tilde{L}_{T_j}$ and $\tilde{L}_{T_j+1}$ would allow us to describe the impact on the wage structure of the arrival of technology j.

The difficulty can be usefully broken down into two problems. First, we do not know, in general, how the menu of technologies effectively used will change between T_j and $T_j + 1$, i.e., how I_{T_j+1} differs from I_{T_j}. All we know is that the former may include j, while the latter will certainly not. Second, even for those i that belong to both I_{T_j} and I_{T_j+1}, we do not know how $L_{T_j+1}^i$ will differ from $L_{T_j}^i$. However, notice that (11) already tells us something useful: *if the wage falls (rises) for one technology that continues to be used after the revolution, it falls (rises) for all such technologies.* Or, if $i, j \in I_{T_j}$, $i, j \in I_{T_j+1}$, and $W_{T_j+1}^i < W_{T_j}^i$, then $W_{T_j+1}^j < W_{T_j}^j$.

I_t and $\tilde{L}_t$ will clearly depend on workers' training decisions. Therefore, we cannot make further progress without imposing some restrictions on the learning-cost distributions $\Phi^i(\sigma_i)$. It turns out that two assumptions about these distributions will be sufficient to generate a number of interesting predictions. The first assumption is a direct extension of the single-revolution model, and consists of assuming that $\Phi^i(\sigma_i)$ is uniform on support $[0, \bar{\sigma}_i]$, where $\{\bar{\sigma}_i\}_{i=0}^{\infty}$ is an exogenously given sequence with $\bar{\sigma}_0 = 0$.[14] For now, assume that $\bar{\sigma}_{i+1} \geq \bar{\sigma}_i$ for all i, but keep in mind that we will relax this monotonicity assumption in the next section. Notice that, using a line of reasoning identical to the one developed in Section I, this first restriction implies that *there will always be at least some workers using the most advanced technology.* This means that $j - 1 \in I_{T_j}$ and $j \in I_{T_j+1}$.

The second restriction is: if one individual occupies the xth percentile of the distribution of the cost of learning technology i, then this individual occupies the xth percentile of all learning-cost distributions. Formally, if one individual has learning costs σ_i and σ_h for technologies i and h, then

$$\frac{\sigma_i}{\bar{\sigma}_i} = \frac{\sigma_h}{\bar{\sigma}_h} = x \qquad \forall i, h.$$

Hence, we can fully identify individuals with the percentile they occupy: e.g., we can refer to the individual above as individual x. While admittedly restrictive, this "rank-preserving" assumption is quite reasonable. Recall the two possible interpretations I have offered for heterogeneous learning costs: differences in cognitive ability and differences in access to credit towards education. Under the former, the assumption says that individuals who are especially "quick" to understand the workings of one type of equipment are also advantaged in picking up the operation of other types of equipment.[15] Under the latter, the rank-preserving property is not only reasonable, but it is probably also the only one to be so.

[13] To obtain equation (11) substitute from the no-arbitrage condition into the definition of K_t. This gives a simple formula for K_t^j/L_t^j that depends on j, K_t, and $\tilde{L}_t$. Now use again the no-arbitrage condition in equation (4) to rewrite W_t^i as a function of K_t^i/L_t^i. Finally, plug in the just-derived formula for K_t^j/L_t^j and simplify.

[14] The requirement $\bar{\sigma}_0 = 0$ insures that workers never have negative net incomes.

[15] This is not to deny that individuals may have different "talents." There are many famous cases of great artists who have been terrible businessmen. But the kind of heterogeneity that is at stake here is not heterogeneity in talents across different spheres of human endeavor (e.g., artistic creativity vs. business acumen), but heterogeneity in ability *within* a particular sphere, namely technical skill.

Individual x will use the technology that affords her the highest net income. Hence, if individual x uses technology i at time t we must have:

$$x(\bar{\sigma}_i - \bar{\sigma}_h) \leq W_t^i - W_t^h \qquad \forall h \in I_t.$$

From this we can immediately derive a *monotonicity property:* if individual x' uses technology i', individual x uses technology i, and $i' > i$, then $x' < x$. This is nothing but saying that the most skilled individuals use the most advanced technologies. In an Appendix that is available upon request I build on this monotonicity property to derive a number of results that partially characterize the equilibrium distribution of employment and net income at a generic date t. The bottom line is as follows.

First, in equilibrium, I_t must be composed of only "contiguous" elements: *if* $h, h' \in I_t$, *and* $h < i < h'$, *then* $i \in I_t$. Following Chari and Hopenhayn (1991) we may term this a "no-hole" condition. It is clearly a consequence of the continuity of the domain of the learning-cost distributions, and it says that if there are workers who find it optimal to use a low-tech technology while some others are using a high-tech technology, there must be at least some workers for whom it is best to adopt an "intermediate-tech" technology. Second, in equilibrium, $\tilde{L}_t$ must satisfy the following pattern:

$$(12) \quad L_t^{g(t)} = \frac{W_t^{g(t)} - W_t^{g(t)-1}}{\bar{\sigma}_{g(t)} - \bar{\sigma}_{g(t)-1}}$$

$$L_t^i = \frac{W_t^i - W_t^{i-1}}{\bar{\sigma}_i - \bar{\sigma}_{i-1}} - \frac{W_t^{i+1} - W_t^i}{\bar{\sigma}_{i+1} - \bar{\sigma}_i}$$

$$\forall i \in I_t, \underline{i} < i < g(t)$$

$$L_t^{\underline{i}} = 1 - \frac{W_t^{\underline{i}+1} - W_t^{\underline{i}}}{\bar{\sigma}_{\underline{i}+1} - \bar{\sigma}_{\underline{i}}},$$

where $\underline{i}$ is defined as the least-advanced technology used at time t: $\underline{i} \in I_t$, $\underline{i} \leq i \; \forall i \in I_t$. The interpretation is that "cost-adjusted wage differentials" (wage differentials among contiguous technologies, divided by learning-cost differentials) effectively partition the unit segment into contiguous segments, each of which represents employment in a given technology.

Now that we have uncovered some of the equilibrium properties of I_t and $\tilde{L}_t$, we can investigate the difference between I_{T_j} and I_{T_j+1} and $\tilde{L}_{T_j}$ and $\tilde{L}_{T_j+1}$. The unpublished Appendix shows that *if* $i \in I_{T_j}$, *and* $i \notin I_{T_j+1}$, *then* $I_{T_j+1} = \{j\}$. If a technology in use before the revolution is abandoned immediately after the revolution, then all technologies in use before the revolution are likewise abandoned. This is a *no partial-upgrading result.* Essentially, it says that a worker choosing i before the revolution will continue to prefer i to any other technology superior to i if that technology was feasible at time T_j. Hence, if i is abandoned, it must be because workers have upgraded "all the way" to j.

Full skill upgrading, or $I_{T_j+1} = \{j\}$ is, however, a possibility. This outcome is entirely equivalent to the "full-adoption" outcome in the one-revolution case: if technology j brings exceptional productivity gains relative to the cost of learning it, all technologies $i < j$ will be immediately abandoned. If this full upgrading does not happen, instead, we are in the "partial-adoption" case we have mostly focused on in the one-revolution version. All the technologies that were in use at time T_j are still in use at time $T_j + 1$, or $I_{T_j} \subset I_{T_j+1}$. This leads us to the key result of this section: *if* $I_{T_j} \subset I_{T_j+1}$ *wage inequality is necessarily greater in period* $T_j + 1$ *than in period* T_j. In fact, wages for workers who do not use technology j in period $T_j + 1$ are lower than the wages they would have received had the revolution not occurred. This conclusion is reached formally in the unpublished Appendix, where it is shown that $W_{T_j+1}^i < W_{T_j}^i$, $\forall i \in I_{T_j}$. The intuition is the same as in the one-revolution case. When the revolution strikes, there will always be a group of workers with low learning cost who will immediately adopt technology j. As these workers must be endowed with capital, there will be a reallocation of capital towards sector j. As rates of return must be equalized across *all* forms of capital, the reallocation of capital is more than proportional to the reallocation of labor towards sector j (because technology j is inherently more productive). This will induce an incipient fall in the capital-labor ratio and hence in the wage *for all workers using preexisting technologies.*[16]

Until now I have been concerned with the initial impact of revolution j. In the partial-adoption case the postrevolutionary dynamics will be described by a generalized version of equation (9). As in the one-revolution case, growth in the total capital stock, together with the no-arbitrage condition, leads capital-labor ratios to increase in concert, leading to widening wage differentials. Hence, the partition of the unit segment implicit in (12) will constantly shift to the right:

$$\frac{W_t^i - W_t^{i-1}}{\bar{\sigma}_i - \bar{\sigma}_{i-1}} > \frac{W_{t-1}^i - W_{t-1}^{i-1}}{\bar{\sigma}_i - \bar{\sigma}_{i-1}}$$

$$T_j + 1 < t \leq T_{j+1}, i - 1 \in I_t.$$

This means that *the postrevolutionary dynamics are characterized by continuous partial skill upgrading:* in every period, some of the workers using technology i would have used technology $i - 1$ had they lived in the previous period. Also, the fraction of the labor force using the lowest technology dwindles over time. Hence, *nonfrontier technologies are gradually abandoned,* and this abandonment proceeds systematically from the lowest technology upward. The first technology to be abandoned is, therefore $\underline{i}'$, and this happens in the first period t such that $(W_t^{\underline{i}'+1} - W_t^{\underline{i}'})/(\bar{\sigma}_{\underline{i}'+1} - \bar{\sigma}_{\underline{i}'}) \geq 1$. For, when this happens, the worker with the highest learning costs, namely $x = 1$, finds it profitable to upgrade her skills by one notch to $\underline{i}' + 1$. Next, technology $\underline{i}' + 1$ is abandoned, and so forth. Notice, however, that this process need not lead to the abandonment of all technologies inferior to j in finite time. Exactly as in the one-revolution case, the economy may reach a steady state in which wages and wage differentials cease to grow, and workers cease to upgrade their skills.

III. De-skilling Technological Revolutions

Until now I have assumed that all revolutions are skill biased. Namely, $\bar{\sigma}_i > \bar{\sigma}_{i-1}$, for all i, so that new technologies involve increasingly large learning costs. As we will see in Section V, however, there are instances in which major technological breakthroughs have been characterized by lower, not higher, learning costs. To refer to such cases, the literature has coined the term "*de-skilling* technological progress." Happily, the model of this paper can be extended to include the study of a de-skilling technological revolution. In order to do this, I must relax the monotonicity assumption on the mapping from i to $\bar{\sigma}_i$.

As before, consider first a simple example in which the economy is using only two technologies when it is hit by a de-skilling revolution. Let us say that the two technologies being used are 0 and 1, and that the new, de-skilling technology is 2 (nothing of substance hinges on this particular choice of labels). Hence, the economy at date T_2 is in the steady state depicted in Figure 1. Finally, assume that $\bar{\sigma}_2 = 0$. It follows from the arguments in Section I that: (i) *both technologies* 0 *and* 1 *are immediately abandoned,* while all the workers, skilled and unskilled, join production method 2; and, obviously, (ii) *inequality declines* as all workers now earn the same wage. In particular, workers at the bottom of the wage distribution earn more than they did before the revolution. As for workers at the top, *the wage earned by the best-paid workers may well decline.*[17] This is because, although they are now

[16] Interestingly, we cannot rule out the possibility of skill "downgrading" as the revolution strikes. Namely, if $\underline{i}'$ ($\underline{i}$) is the "lowest" technology used in $T_j + 1$ (T_j), the only restriction is $\underline{i} \geq \underline{i}'$. The intuition is that, as absolute wages fall for less-than-frontier technologies, so do wage differentials. Faced with a fall in the wage differential between two technologies, and given the constancy of learning-cost differentials, some workers in some technologies may find that the gain is no longer worth the cost. In other words, the technological revolution spreads the menu of technologies used not only "up," with the addition of j, but also "down," with the recuperation of previously abandoned techniques. In studying the transitional dynamics following period $T_j + 1$, however, we will soon see that while the former addition is long-lived, the latter is typically short-lived.

[17] To see this notice that $W_{T_2+1}^2 = A^2(1 - \alpha)(K_{T_2})^\alpha$ from (4) while $W_{T_2}^1 = A^{[1/(1-\alpha)]}(1 - \alpha)(K_{T_2})^\alpha/[1 + L_{T_2}^1(A^{[1/(1-\alpha)]} - 1)]^\alpha$ from (11). The difference between these two cannot be signed in the allowed region of the parameter space.

using more productive machines (of type 2 instead of type 1), they have to share them with the previously unskilled. This dilutes their capital-labor ratio. The net effect is ambiguous.

If $0 < \bar{\sigma}_2 \leq \bar{\sigma}_1$, the picture is only slightly modified. Now, only 1 will certainly be immediately and permanently abandoned, while 0 may remain in use if for some of the least-skilled workers 2 is still too costly. However, it is certainly the case that $L^2_{T_2+1} > L^1_{T_2}$, i.e., the proportion of the labor force that finds it profitable to use the high-tech technology (now 2) has increased. This is because the new machines are *both* more productive *and* easier to use. In this case, the effects on the wage structure are as follows: W^0 falls, while W^2 may be either less or more than W^1.[18] The workers left behind in technology 0 suffer a further decline in their capital-labor ratio, and hence in their earnings. The workers who would previously have been in technology 1 will all be in the more productive technology 2 after the revolution, but—due to $L^2_{T_2+1} > L^1_{T_2}$—they will also suffer a dilution of their capital endowment. Even if the net effect is $W^2_{T_2+1} > W^1_{T_2}$, however, it is important to realize that the fall in W^0 does not mean that inequality increases. For, there is a (potentially large) fraction of workers now utilizing machines of type 2, whose correspondents in the previous generation were using machines of type 0. For this segment of the labor force wages increase. Inequality may therefore fall, and different measures of inequality may move in opposite directions.

Now for the general case. With $\{\bar{\sigma}_i\}_{i=0}^{\infty}$ nonmonotonic, the results and the proofs of the previous section require some notational amendments, but they all go through in substance. For example, the "no-hole" condition now becomes: if $h, h' \in I_t$, $h < i < h'$, *and* $\bar{\sigma}_i < \bar{\sigma}_{h'}$, then $i \in I_t$. In other words, there will be holes corresponding to technologies (and only to technologies) that are *less* productive and *more* costly to learn than other available technologies. Because these technologies are abandoned forever, however, the existence of this type of holes does not affect any of the results derived elsewhere, although the index i must be reinterpreted somewhat. Imagine now that the economy is in steady state at time T_j, and this steady state is described by generic I_{T_j} and $\tilde{L}_{T_j}$. Call $\bar{i}$ the greatest element of I_{T_j}. Revolution j strikes with $\bar{\sigma}_j \leq \bar{\sigma}_{\bar{i}}$. The results derived above are sufficient guidance to identify the impact effect. First, all technologies h such that $\bar{\sigma}_j \leq \bar{\sigma}_h$ are necessarily immediately and permanently abandoned. The mass of workers using technology j after the revolution is larger than the mass of workers in all the now-abandoned technologies before the revolution, and may well include the entire labor force. Wage inequality may either increase or decline. The postrevolutionary dynamics are exactly as described in the previous section.

IV. Inequality Facts, Old and New

A. Existing Evidence

The model of Sections I and II offered an explanation for the recent increase in U.S. wage inequality. The existing empirical literature contains many other findings that accord well with the model, and this subsection briefly reviews them. The next subsection documents some new empirical findings that are also in broad agreement with the message of this paper. For simplicity in linking the empirical review to the theory I will refer to the one-revolution case of Section I.

Evidence from Wage Regressions.—First, there has been a large increase in the returns to schooling (e.g., John Bound and George Johnson [1992] for the United States; John Schmitt [1995] for the United Kingdom).

[18] We have:

$$W^0_{T_2+1} = \frac{(1-\alpha)(K_{T_2})^\alpha}{[1 + L^2_{T_2+1}(A^{[2/(1-\alpha)]} - 1)]^\alpha}$$

$$< \frac{(1-\alpha)(K_{T_2})^\alpha}{[1 + L^1_{T_2}(A^{[1/(1-\alpha)]} - 1)]^\alpha} = W^0_{T_2}$$

and

$$W^2_{T_2+1} = \frac{A^{[2/(1-\alpha)]}(1-\alpha)(K_{T_2})^\alpha}{[1 + L^2_{T_2+1}(A^{[2/(1-\alpha)]} - 1)]^\alpha}$$

$$\lesseqgtr \frac{A^{[1/(1-\alpha)]}(1-\alpha)(K_{T_2})^\alpha}{[1 + L^1_{T_2}(A^{[1/(1-\alpha)]} - 1)]^\alpha} = W^1_{T_2}.$$

Imagine that schooling improves or is correlated with learning ability. Then education proxies (inversely) for the learning cost σ: more educated (low σ) workers will be quicker to adopt the new technology, and will therefore receive higher wages. Second, there has been an increase in residual wage inequality, i.e., in the variance of the residuals in log-wage regressions. This is often interpreted as evidence of increased returns to unobserved abilities. In the context of the model the relevant unobserved abilities are unobserved learning skills, i.e., the component of learning ability that is uncorrelated with education, experience, etc. Hence, the increased residual inequality may capture increased returns to low σ's. Third, individuals who use a computer at work earn higher wages (Alan B. Krueger, 1993). If computers are type-1 machines, this is also consistent with the model.[19]

Evidence from Plant-Level Data.—(i) Average wages and workers' education are higher in manufacturing plants with high measures of R&D and adoption of new technologies (Timothy Dunne and James A. Schmitz, 1995; Mark Doms et al., 1997): type-1 plants pay high wages and have low-σ workers. (ii) Plants installing high-tech equipment have relatively skilled labor forces both before and after the adoption (Dunne et al., 1997). This suggests that plant managers recognize that they need a skilled labor force to work with the new machinery. Hence, plants where the workforce is particularly unskilled will continue to use the old technology. (iii) There is a significantly positive correlation between indices of adoption of advanced technologies and investment (as a percent of the capital stock) (Doms et al., 1997; Dunne et al., 1997). Hence, plants that are using more advanced technologies are also investing more, just as the model predicts. (iv) There has been an ongoing upgrading of the skill level of the labor force at the plant level. Skill upgrading is positively correlated across plants with the capital-output ratio (Doms et al., 1997; Dunne et al., 1997). (v) There has been a very substantial increase in between-plant average wage inequality, and differences among plants in the size of the labor force explain a very large fraction of this increase (Steve J. Davis and John Haltiwanger, 1991). But plant size is almost certainly highly positively correlated with capital intensity.[20] Hence, this finding suggests that cross-plant heterogeneity in capital-labor ratios may have played an important role in generating increased cross-plant wage dispersion.

Evidence from Industry-Level Data.—When the average age of an industry's capital stock falls, the wage share of highly educated workers increases (Ann P. Bartel and Frank R. Lichtenberg, 1987). Recently, returns to education have increased the most in those industries in which investment in new technology has been more rapid (Steven G. Allen, 1996). These findings support the view that—in the recent period—skills have played a driving role in technology adoption. Another relevant fact is that within-industry skill upgrading—which accounts for most of the economywide increase in the relative demand for nonproduction workers—is positively correlated with increases in the industry average capital-output ratio (Eli Berman et al. [1994] for the United States; Stephen Machin [1996] for the United Kingdom; Berman et al. [1997] for OECD countries).

B. *The Empirical Dispersion of Capital-Labor Ratios*

The large increase in wage inequality documented in the literature invites the question of what happened to the distribution of capital across workers. I now provide some preliminary evidence on this using industry-level data from the NBER Productivity Database. The database covers 450 manufacturing industries from 1958 to 1991, and includes information on payroll, employment of production and

[19] However, John E. DiNardo and Jörn-Steffen Pischke (1997) attribute the positive coefficient on computer usage to reverse causation.

[20] Across four-digit manufacturing industries the size of the average plant is highly correlated with average capital intensity. I thank Steve Davis for sharing this calculation with me.

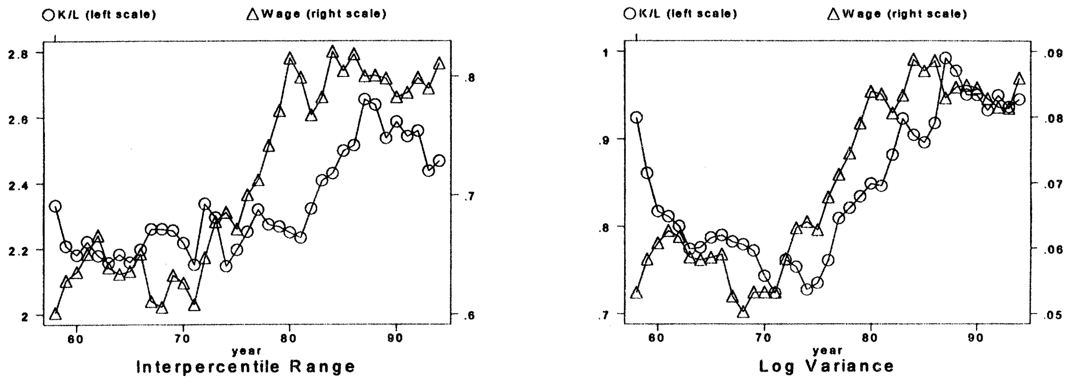

FIGURE 3. INTERINDUSTRY DISPERSION OF EQUIPMENT PER WORKER AND PAY PER WORKER

nonproduction workers, and real stocks of equipment and structures.[21] Dividing the equipment-stock figure by total employment I obtain a simple measure of capital per worker in each industry and time period.

In Figure 3, circles identify plots of the time paths of the dispersion of the equipment/employment ratio across industries. The left panel measures inequality by the difference between the logarithms of the 90th and the 10th percentile of the distribution (interpercentile range). The right panel shows the evolution over time of the variance of the logarithm of the equipment/employment ratio (log variance). As measured by the interpercentile range, the interindustry dispersion of the equipment/employment ratio appears fairly stable throughout the 1960's and most of the 1970's, but increases sharply—about 18 percent—in the 1980's. Unlike the interpercentile range, the log variance declines throughout the 1960's. However, its time path in the 1970's and 1980's closely mimics the behavior of the interpercentile range. In fact, the percentage increase in inequality of capital intensity between 1975 and 1990 registered by the log variance—about 31 percent—is even greater than the one implied by the interpercentile range. If, instead of using equipment, I measure capital by the sum of equipment and structures, I obtain qualitatively identical plots. When I weight each industry's capital-labor ratio by the time average of its share in total employment (fixed weights), I still obtain the same picture. However, an important caveat emerges when I weight by the current industry share in total employment (variable weights): in this case, the interpercentile range of the interindustry dispersion of capital per worker only increases by 8.4 percent between 1975 and 1990, while the log variance actually *falls* slightly (1.9 percent). I return below to this inconsistency between the behavior of variable weighted and fixed weighted or unweighted measures of inequality.

The new fact documented in Figure 3 is that the (unweighted) cross-industry dispersion of capital-labor ratios has increased rapidly and substantially in the recent past. The question, however, is whether the two phenomena—increased inequality in wages, and increased inequality in capital-labor ratios—are related to each other. Figure 3 also shows measures of interindustry wage inequality (denoted by triangles), where industry wages are computed as pay per worker (divided by the CPI). The comparison between interindustry inequality in wages and in capital intensity is interesting but inconclusive. It seems that the two series track each other very closely, but with a time lag. Specifically, the upward trend in wage inequality appears to start around 1971, while the trend towards increased inequality of

[21] See Eric Bartelsman and Wayne Gray (1994) for a detailed description of the data set.

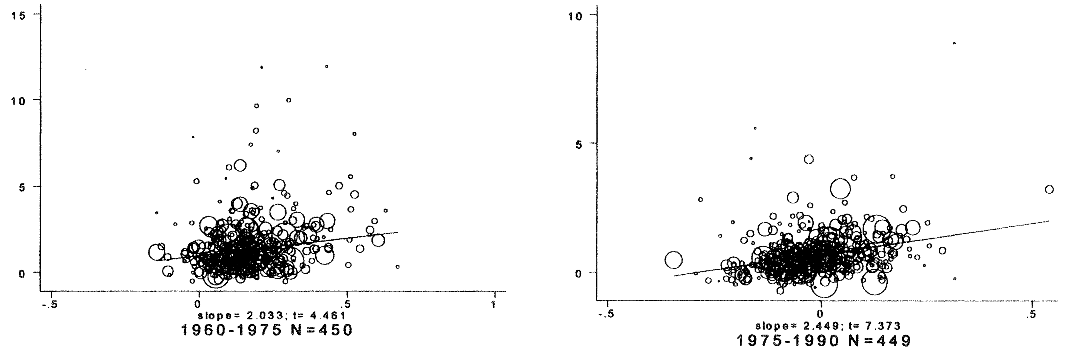

FIGURE 4. CHANGES IN CAPITAL PER WORKER VS. CHANGES IN PAY PER WORKER

capital-labor ratios starts in 1974 at the earliest.

To further explore the relationship between wages and capital intensity, Figure 4 plots (percent) changes in the equipment-employment ratio against (percent) changes in wages. In the left panel the changes are between 1960 and 1975. In the right panel they are between 1975 and 1990. Figure 4 shows a strong and positive relationship between changes in wages and changes in capital-labor ratios.[22] The strongest evidence that increased capital-labor ratio inequality is associated to the increase in wage inequality is provided in Figure 5, which shows scatterplots of the proportional change in the equipment-employment ratio against *initial* pay per worker for the pre- and post-1975 subperiods. The correlation between initial wages and successive changes in the capital-labor ratio turns out to be negative in 1960–1975, and significantly positive in 1975–1980. Hence, capital has flowed towards industries that had high wages to start with. As a consequence, wages in these industries have increased, and so has wage inequality.

It is also important to establish a connection between these changes in the distribution of capital and industry-level skill intensity, as measured by the share of nonproduction workers in total employment.[23] Figure 6 shows that industries with above-average increases in their capital-labor ratio also had above-average skill upgrading—increases in the share of skilled workers in total employment—between 1975 and 1990. This finding was already (implicitly) featured in Berman et al. (1994), who showed that skill upgrading positively correlates with increases in the capital-*output* ratio. It is remarkable that, instead, changes in the share of nonproduction workers were *negatively* correlated with changes in capital intensity between 1960 and 1975. Hence, capital-skill complementarity per se does not automatically generate a positive relation between changes in capital-labor ratios and skill upgrading. This positive relation is a feature of the information-technology period, but not of the previous 15 years. Figure 7 depicts changes in capital-labor ratios against initial skill intensity: initially skill-intensive industries—which might plausibly be expected to be at the forefront of the tech-

[22] Each circle represents an industry and the size of the circle is proportional to the time average of the corresponding industry's share in total employment. For each plot, I draw a (weighted) regression line and report slope coefficient and standard error. To improve clarity, outlying industries or industries with missing values have been omitted. Nowhere in this section are the results sensitive to inclusion or exclusion of outliers.

[23] The identification of nonproduction workers with skilled labor, and production workers with unskilled labor—problematic in principle—has been convincingly defended on empirical grounds by, e.g., Dunne et al. (1997).

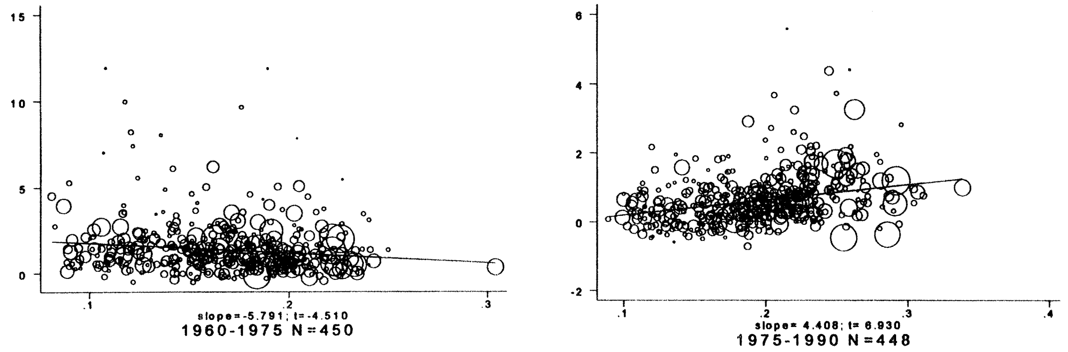

FIGURE 5. CHANGES IN CAPITAL PER WORKER VS. INITIAL IN PAY PER WORKER

nological revolution—underwent disproportionate increases in their capital endowment per worker in the postrevolutionary period. In contrast, in the 1960–1975 period, the largest increases in capital per worker went to unskilled-intensive industries.

Is the model of Section I consistent with these findings? Consider the simple model of one technological revolution, but imagine that there are now several industries. Each industry takes part in the revolution and engages in some degree of adoption of equipment and labor of type 1. Hence, within-industry wage inequality increases as each industry has both workers of type 0 and workers of type 1. Imagine, however, that industries differ in the degree to which they adopt the new machines. In other words, for some industries postrevolution output is largely created by type-1 technology, while other industries lag behind and obtain most of their output from type-0 capital and labor.[24] Then, an increase in inequality of capital-labor ratios and wages among workers within industries will also generate an increase in inequality of average capital-labor ratios and average wages between industries. Industries with high rates of adoption will have a high incidence of highly paid workers, who become endowed with higher capital-labor ratios. Industries with low adoption rates will have a disproportionate amount of high learning-cost workers, who experience declining capital-labor ratios and wages.[25]

Further predictions can be obtained by assuming that industry shares in GDP are constant, perhaps because consumer utility functions are Cobb-Douglas in the various industry outputs. Then, total employment will shrink in industries with high adoption of the new technology, and expand in industries with low adoption. Also, there will be a negative cross-industry correlation between changes in total capital and changes in total employment.[26] The intuition is that high-adopting industries are now more productive, and therefore need fewer workers to generate the same amount of output. Instead, low-adopting industries are less productive—because they have less capital—and therefore need more labor. This behavior of the employment shares can explain the puzzling finding—reported above—that current-employment weighted measures of inequality in the capital-labor ratio did not increase after 1975, while the unweighted measures did. To see this, imagine that there are only two industries, x and z, and that we measure inequality by the log difference of the capital-labor ratios, $\log(k_z) -$

[24] That the extent to which the new technology is adopted varies substantially across industries is documented by Berman et al. (1994) and Allen (1996).

[25] A note extending the one-sector model of Section I to a two-sector model is available upon request.

[26] This is borne out by the data: the correlation between log changes in equipment and log changes in employment is −0.41.

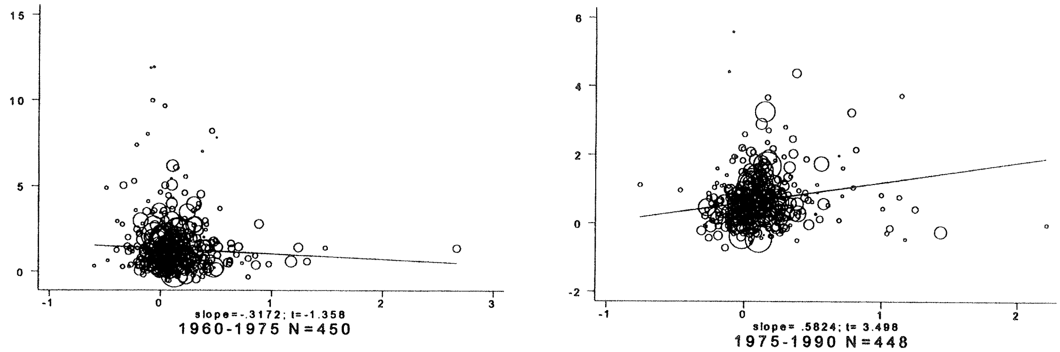

FIGURE 6. CHANGES IN CAPITAL PER WORKER VS. CHANGES IN NONPRODUCTION-WORKER SHARE

$\log(k_x)$. The theory predicts that k_z increases and k_x falls, so unweighted inequality increases. However, weighted inequality is $\log(k_z)L_z - \log(k_x)L_x$, where L_i is total employment in industry $i = x, z$. The theory predicts that L_z falls while L_x increases. Hence, while the unweighted variance always increases, the weighted variance may well decline. A numerical simulation available from the author shows that this intuition carries over to the log variance.[27]

V. Technological Revolutions of the Past Two Centuries

Does the framework developed in this paper have some usefulness in interpreting facts beyond the last quarter century? The prototypical technological revolution of the modern era is the steam engine (1765). As a new tool for the generation of power, the steam engine is likely to have been skill biased, as it was both mechanically more complex, and involved more sophisticated principles than the sources used before its advent (humans, animals, wind, and water). More broadly, in several manufacturing industries (most notably textiles, metals, food, drink, and tobacco) the application of steam power went hand in hand with the introduction of new types of machines (Nick von Tunzelmann, 1986), some of which may have required extra investments in skills on the part of the workers. For example, some cotton was spun by steam-powered mules already in the 1790's, and woven by power looms circa 1810–1820 (Joel Mokyr, 1990). In the case of power looms, the case for a skill-biased component is reinforced by the fact that a large fraction of the hand-loom weavers displaced by power looms became common laborers, instead of joining power-loom operations (Peter H. Lindert, 1994). And this despite the fact that cotton production was booming, and certainly not afflicted by excess supply of workers (Nicholas F. R. Crafts, 1985).

Several pieces of evidence are consistent with the theoretical model. First, there was a protracted period of coexistence of the old and the new technologies. In the generation of power, water remained extremely important alongside steam until at least the 1850's (Mokyr, 1990). Similarly, the hand loom continued to be used in cottage settings for about 40 years after the introduction of the power loom (Antonio Ciccone, 1997). Second, the capital-labor ratio for power-loom operators was at least three times as large as the corre-

[27] Pushing the multisector extension further, it is possible to make sense also of the positive correlation between initial wages (and skill intensity) and changes in capital-labor ratios. With multiple revolutions, high-wage, high capital-labor ratio sectors are those with the highest concentration of low learning-cost workers. Hence, they will be the first to embrace a new revolution and, as a result, a further reallocation of capital in favor of these high-wage sectors will take place. In the model, as in the data, capital flows from low-wage to high-wage industries, thereby further increasing wage inequality.

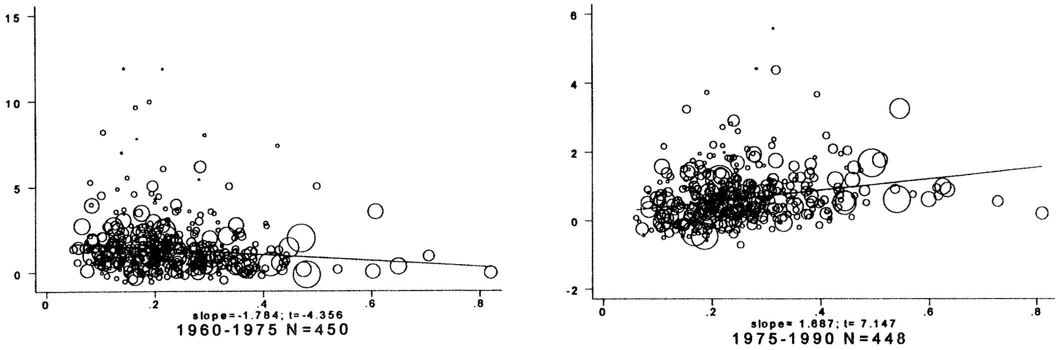

FIGURE 7. CHANGES IN CAPITAL PER WORKER VS. INITIAL NONPRODUCTION-WORKER SHARE

sponding ratio for hand-loom operators (even without counting the steam engine as part of the equipment) (Ciccone, 1997). Third, it seems that between 1810 and 1830 the skill premium increased, although this increase follows a decline in the previous two decades (Lindert, 1994).[28] Fourth, at least some unskilled-worker wages declined. Such is the case of the hand-loom weavers, whose "earning power was ... destroyed ... , after 1822, by new competition from the power looms of the factories" (Lindert, 1994 p. 372).

In the 1880's dynamo technology was ready for introduction in manufacturing, and this led to the electrification of industry (Mokyr, 1990). Naturally, electrification meant the appearance on the factory floor of a whole wave of new gadgets and machines ("conveyors, traveling cranes, jitneys, carriers, industrial trucks, and other handling devices"), and the disappearance of others (Claudia Goldin and Lawrence F. Katz, 1998 p. 713). Did the new electrical machinery require a greater learning investment? Several facts are in accordance with this view. First, it took dynamos a long time to replace steam engines: penetration of electricity as a mechanical drive for manufacturing equipment was 5 percent in 1899, 25 percent in 1909, 53 percent in 1919, and 78 percent in 1929 (Paul A. David, 1991). Second, between 1909 and 1929 industries that used more electricity (as opposed to other sources of power) paid higher wages, employed a more educated workforce, and had a higher capital-labor ratio (Goldin and Katz, 1998). All these facts match the predictions of this paper's model under the assumption that the Dynamo Revolution was skill biased.

Since Charlie Chaplin's *Modern Times,* the assembly line is the epitome of de-skilling technological change. While assembly lines are a late nineteenth-century invention (Mokyr, 1990), the "paradigm-setting" experience is that of Henry Ford's Highland Park facility, in 1913. Until 1910 or so, car making had been the domain of highly skilled artisans, fitting the parts and assembling the cars in a process in which division of labor was minimal. With the assembly line "every*thing* was put into motion and every *man* brought to a halt." Thousands of unskilled laborers were soon collecting an unbelievable 5 dollars a day, while artisanal production of cars was all but wiped out in the space of a few years (David A. Hounshell, 1984 p. 244). Clearly, this description fits well with the representation of a de-skilling technological revolution given in Section III. Wages for the unskilled rise while those for the skilled decline. The fact that there is (almost) immediate complete abandonment of the old technology is also consistent with the theory, and is one of the basic differences between a skill-biased (that admits long overlaps between old and new technology) and a de-skilling (immediate abandonment) revolution.

[28] The turn-of-the-century decline may perhaps be attributable to the de-skilling nature of some of the many other innovations taking place in the textile industry in the second half of the eighteenth century.

VI. Summary

This paper contributes to our understanding of the Information-Technology Revolution in three main respects. First, it clarifies the mechanism by which, on impact, the revolution generates absolute gains for those individuals with high cognitive ability (or no credit constraints), and absolute losses for those with high costs of learning. Second, it discusses the potential future evolution of the wage structure. Ever-widening wage differentials and, potentially, learning externalities will draw an increasing number of workers in the skilled pool, thereby reducing inequality. However, decreasing marginal returns to capital may lead the economy towards a steady state in which the labor market remains split between skilled and unskilled workers. The third contribution is empirical. The main finding is that the recent increase in wage inequality is associated with increased inequality in capital-labor ratios. While of independent interest, this fact conforms well with the predictions of my model. An overview of other episodes in the history of technology indicates that the model may have interpretative power for other technological revolutions, including de-skilling ones.

APPENDIX A: EQUILIBRIUM OF PERIOD $T + 1$

It is easiest to organize the search for an equilibrium for period $T + 1$ around the three conceivable cases given by $R^0_{T+1} \gtreqless R^1_{T+1}$.

(1) $R^1_{T+1} < R^0_{T+1}$.—This candidate equilibrium is easily ruled out. The condition implies $K^1_{T+1} = 0$. Hence, $W^1_{T+1} = 0$, $R^1_{T+1} = 0$. Also, $K^0_{T+1} = \bar{K}$, $L^0_{T+1} = 1$, $W^0_{T+1} = \bar{W}$, and $R^0_{T+1} = \bar{R}$. Call $\Psi(\cdot)$ the inverse function of $\Phi^1(\cdot)$. Now consider a firm in sector 1 that rents $x\bar{K}$ units of capital at the rental price $\bar{R}$ and x units of labor at the wage $\bar{W} + \Psi(x)$. Given the assumptions on factor supplies it will always be able to do so: at the offered prices investors will be indifferent between renting to this firm or to a firm in sector 0, while all workers with learning cost less than $\Psi(x)$ will strictly prefer to accept this firm's offer. It is now easy to see that this firm can always choose x so as to make a profit, thereby contradicting the 0-profit condition. To see this, just notice that profits for this firm are:

$$A x \bar{K}^{\alpha} - x[\bar{W} + \Psi(x)] - x\overline{KR}$$
$$= x\{[\bar{K}^{\alpha} - \bar{W} - \overline{KR}] + (A-1)\bar{K} - x\bar{\sigma}\}.$$

The claim follows immediately after observing that the expression in square brackets is 0.

(2) $R^1_{T+1} = R^0_{T+1}$.—The first thing to observe is that interest-rate equalization necessarily requires strictly positive employments of both factors in both sectors. If labor or capital were 0 in either sector 0 or sector 1, the factor-pricing equations would lead to an immediate contradiction of the interest parity. This means that we can concentrate on the "interior" portion of the labor-1 supply function (7): $L^1_{T+1} = (W^1_{T+1} - W^0_{T+1})/\bar{\sigma}$. This function, the interest parity, the feasibility constraints $L^0_{T+1} = 1 - L^1_{T+1}$ and $K^0_{T+1} = \bar{K} - K^1_{T+1}$, and the four factor-pricing equations, constitute a system of eight equations in eight unknowns. This system could in principle be solved for W^1_{T+1}, R^1_{T+1}, L^1_{T+1}, K^1_{T+1}, W^0_{T+1}, R^0_{T+1}, L^0_{T+1}, and K^0_{T+1}, all as functions of the exogenous parameters and the state variable $\bar{K}$. Some substitutions allow to reduce the system into two equations in the two unknowns L^1_{T+1} and K^1_{T+1}:

$$L^1_{T+1} = \left[\frac{(1-\alpha)(A - A^{[\alpha/(\alpha-1)]})}{\bar{\sigma}}\right]^{[1/(1+\alpha)]} (K^1_{T+1})^{[\alpha/(1+\alpha)]}$$

and

$$L^1_{T+1} = \frac{A^{[1/(\alpha-1)]}K^1_{T+1}}{\bar{K} - (1 - A^{[1/(\alpha-1)]})K^1_{T+1}}.$$

By examining the properties of these two curves in (K^1_{T+1}, L^1_{T+1}) space, one quickly reaches the conclusion that the system always admits a solution for any positive $\bar{K}$, and that in the solution both L^1_{T+1} and K^1_{T+1} are monotonically decreasing in $\bar{K}$. This solution represents an equilibrium for period $T + 1$ if and only if it involves $L^1_{T+1} < 1$ and $K^1_{T+1} < \bar{K}$. Now notice that $(L^1_{T+1} = 1, K^1_{T+1} = \bar{K})$ is the solution whenever $\bar{K} = \tilde{K}$, where

$$\tilde{K} \equiv \left[\frac{\bar{\sigma}}{(1-\alpha)(A - A^{[\alpha/(\alpha-1)]})}\right]^{1/\alpha}.$$

Because the solution is unique and monotonically decreasing in $\bar{K}$, the desired condition is $\bar{K} < \tilde{K}$. Once the existence of a solution $L^1_{T+1} \in (0, 1)$ has been established, the equations of the original system readily deliver the expressions in (8).

(3) $R^1_{T+1} > R^0_{T+1}$.—The interest rate inequality implies, from the supply correspondences, $K^1_{T+1} = \bar{K}$, $K^0_{T+1} = 0$ and hence, via the factor-pricing equations, $W^0_{T+1} = 0$. Also, L^0_{T+1} must be 0, otherwise R^0_{T+1} would be infinity. In turn this implies $L^1_{T+1} = 1$ and $R^0_{T+1} = 0$. Finally, the interest rate is $R^1_{T+1} = \alpha A\bar{K}^{\alpha-1}$ and the wage is $W^1_{T+1} = (1-\alpha)A\bar{K}^{\alpha}$. These prices support an equilibrium only if firms cannot make profits by luring workers "back" to sector 0. Imagine a firm approaching the worker with highest learning cost with a proposal to work for wage $W^1_{T+1} - \bar{\sigma}$ on $x\bar{K}$ units of type-0 capital. To the owners of capital the firm pays interest R^1_{T+1}. Being indifferent, both worker and capital owners will accept this arrangement, which will then provide to the firm profits:

$$(x\bar{K})^{\alpha} - (1-\alpha)A\bar{K}^{\alpha} - x\alpha A\bar{K}^{\alpha} + \bar{\sigma}.$$

Maximizing this with respect to x one obtains $x = A^{[1/(\alpha-1)]}$. For our candidate equilibrium to be admissible it is necessary that maximum profits from the above-described deviation be negative. Plugging the optimal value of x in the profit function this no-deviation condition turns out to be the exact opposite of the condition supporting the equilibrium in point (2), namely $\bar{K} \geq \tilde{K}$. Notice that whenever this condition is satisfied, we also have that $(1-\alpha)A\bar{K}^{\alpha} \geq \bar{\sigma}$ so that all workers do indeed find it profitable to work with tools of type 1.

Hence, whenever $\bar{K} < \tilde{K}$, we have a unique equilibrium with partial adoption. Whenever $\bar{K} \geq \tilde{K}$, we have a unique equilibrium with full adoption. Notice that $\tilde{K}$ is decreasing in A and increasing in $\bar{\sigma}$. In the text I have assumed that A is sufficiently low relative to $\bar{\sigma}$ that $\bar{K} < \tilde{K}$ holds.

APPENDIX B: PROPERTIES OF FUNCTIONS $h(\cdot)$, $q(\cdot)$, AND $\hat{q}(\cdot)$

As discussed in point (2) of Appendix A, in periods of partial adoption the system reduces to two equations in L^1 and K^1 (I am dropping time subscripts). These can be further reduced to the expression:

$$\text{(B1)} \qquad (1-b)(L^1)^{[(1+\alpha)/\alpha]} + b(L^1)^{(1/\alpha)} - cK = 0,$$

where $b \equiv A^{[1/(\alpha-1)]} < 1$, and $c = 1/\tilde{K}$. Equation (B1) implicitly defines $h(K)$. We know from Appendix A that $h(0) = 0$, that $h(K) > 0$ for $K > 0$, and that $h'(K_t) > 0$. Using the implicit-function theorem we find that

$$h'(K) = c/\{(1-b)\frac{1+\alpha}{\alpha}[h(K)]^{(1/\alpha)} + b\frac{1}{\alpha}[h(K)]^{[(1-\alpha)/\alpha]}\}$$

which shows that $h'(0) = \infty$ and $h'(\infty) = 0$. By reapplying the implicit-function theorem, together with the chain rule, we get (after collecting terms):

$$h''(K) = -\{c^2\frac{1}{\alpha}[(1-b)\frac{1+\alpha}{\alpha}[h(K)]^{(1/\alpha)} + b\frac{1-\alpha}{\alpha}[h(K)]^{[(1-\alpha)/\alpha]}]\} / \{h(K)[(1-b)\frac{1+\alpha}{\alpha}[h(K)]^{(1/\alpha)} + b\frac{1}{\alpha}[h(K)]^{[(1-\alpha)/\alpha]}]^3\}$$

which is negative for $K > 0$.

We can now turn to $q(K_t)$ (I return to subscripts). Using the uniform-distribution assumption, the L^1_t supply function, the factor-pricing equations, and the interest parity condition, this can be rewritten as:

$$q(K_t) = \frac{\beta}{1+\beta}\bar{\sigma}\left[\frac{1}{2}[h(K_t)]^2 + \frac{1}{A^{[1/(1-\alpha)]}-1}h(K_t)\right].$$

Using the results developed above for $h(\cdot)$ it is clear that this is an increasing function, passing through the origin with infinite slope and tending towards infinity with a slope of zero. Differentiating twice we get:

$$q''(K_t) = \frac{\beta}{1+\beta}\bar{\sigma}\left\{[h'(K_t)]^2 + \left[h(K_t) + \frac{1}{A^{[1/(1-\alpha)]}-1}\right][h''(K_t)]\right\}.$$

Substituting for $h'(K_t)$ and $h''(K_t)$ from above it is fairly straightforward to show that this expression is negative, so that $q(K_t)$ is concave.

To see that $\hat{q}(K) < q(K)$ for $K < \tilde{K}$ suppose, by contradiction, that there exists $\underline{K} < \tilde{K}$ such that $\hat{q}(\underline{K}) > q(\underline{K})$. This implies

$$(1-\alpha)A\underline{K}^{\alpha} - \frac{1}{2}\bar{\sigma} > \bar{\sigma}\left[\frac{1}{2}[L^1(\underline{K})]^2 + \frac{1}{A^{[1/(1-\alpha)]}-1}[L^1(\underline{K})]\right].$$

Using the labor-1 supply function this can be rewritten:

$$\text{(B2)}\quad (1-\alpha)A\underline{K}^{\alpha} - \frac{1}{2}\bar{\sigma} - \frac{1}{2}\frac{[W^1(\underline{K}) - W^0(\underline{K})]^2}{\bar{\sigma}} + \frac{1}{A^{[1/(1-\alpha)]}-1}[W^1(\underline{K}) - W^0(\underline{K})] > 0.$$

Now imagine that, when the economy reaches the state $\underline{K}$, a company offers the following deal. To capitalists, it offers to rent all their capital at the rate $\underline{R} = \alpha A\underline{K}^{\alpha-1}$. Capitalists will go for it since $\underline{K} < [K^1(\underline{K})]/[L^1(\underline{K})]$ and hence $\underline{R} > R^1(\underline{K})$. To workers whose learning cost is less than $[W^1(\underline{K}) - W^0(\underline{K})]$, it offers wage $W^1(\underline{K})$. These workers will accept out of indifference. Finally, to workers whose σ is above $[W^1(\underline{K}) - W^0(\underline{K})]$, the company offers $W^0(\underline{K}) + \sigma$. These workers are also indifferent and therefore will accept. The profits from this operation are:

$$A\underline{K}^{\alpha} - \underline{RK} - W^1\frac{W^1 - W^0}{\bar{\sigma}} - W^0\left(1 - \frac{W^1 - W^0}{\bar{\sigma}}\right) - \int_{W^1 - W^0}^{\bar{\sigma}} \frac{\sigma}{\bar{\sigma}}\,d\sigma.$$

With a few substitutions this expression is easily seen to coincide with the left side of (B2), and therefore to be positive if $\hat{q}(\underline{K}) > q(\underline{K})$. This is a contradiction.

Appendix C: Slow Depreciation

In the main text I have assumed that equipment fully depreciates within the lifetime of one generation. If capital of type 0 can costlessly be transformed in capital of type 1, and vice versa, slow depreciation can be introduced without modifying any of the formulas in the paper. Multiperiod depreciation generates potential problems, therefore, only if there is a cost associated with the conversion of capital from one type to another. To fix ideas, let us consider the extreme case in which this cost is infinity. Call δ the rate of depreciation of capital per period. At the time of the technological revolution there is a supply of $(1-\delta)\bar{K}$ units of type-0 capital on the used-capital market. In equilibrium these must find a buyer. Define $\tilde{K}^0_{T+1}$ the equilibrium period $T+1$ level of type-0 capital in the one-period depreciation case studied in the main text. If $\tilde{K}^0_{T+1} > (1-\delta)\bar{K}$, the rate of depreciation is, again, immaterial even for the formulas of the paper: on the used-capital market supply is fully absorbed at the price (in units of the consumption good) of one. If $\tilde{K}^0_{T+1} < (1-\delta)\bar{K}$, instead, the existence of an irreversible stock of type-0 capital becomes a "binding constraint" on the ability of the economy to reallocate resources towards sector 1. We will therefore have $K^0_{T+1} = (1-\delta)\bar{K}$, and $K^1_{T+1} = [1 - p_T(1-\delta)]\bar{K}$, where p_t is the price of used (type-0) capital in period t. A formal characterization of the equilibrium in this case becomes difficult (too many state variables). However, it seems possible that such equilibrium may *not* involve an absolute fall in the wage of workers remaining in sector 0. To see this, notice that the no-arbitrage condition becomes:

$$\alpha A\left(\frac{K^1_{T+1}}{L^1_{T+1}}\right)^{\alpha-1} + (1-\delta)$$
$$= \frac{1}{p_T}\left[\alpha\left(\frac{K^0_{T+1}}{L^0_{T+1}}\right)^{\alpha-1} + p_{T+1}(1-\delta)\right].$$

Depending on the time path of p_t, this condition does not necessarily require a higher capital-labor ratio in sector 1 relative to sector 0. Of course, even in cases in which the technological revolution does not cause an *absolute* wage loss for the workers who stay in sector 0, it will still be true that wage inequality goes up (skilled workers would not move to sector 1 otherwise) and hence absolute skilled-worker wages will always increase. Postrevolutionary dynamics will likewise be similar to those described in the text.

REFERENCES

Acemoglu, Daron. "Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality." *Quarterly Journal of Economics,* November 1998, *113*(4), pp. 1055–89.

Allen, Steven G. "Technology and the Wage Structure." National Bureau of Economic Research (Cambridge, MA) Working Paper No. 5534, 1996.

Bartel, Ann P. and Lichtenberg, Frank R. "The Comparative Advantage of Educated Workers in Implementing New Technology." *Review of Economics and Statistics,* February 1987, *69*(1), pp. 1–11.

Bartelsman, Eric and Gray, Wayne. "NBER Manufacturing Productivity Database." Unpublished manuscript, National Bureau of Economic Research, 1994.

Berman, Eli; Bound, John and Griliches, Zvi. "Changes in the Demand for Skilled Labor within U.S. Manufacturing: Evidence from the Annual Survey of Manufacturing." *Quarterly Journal of Economics,* May 1994, *109*(2), pp. 367–97.

Berman, Eli; Bound, John and Machin, Stephen. "Implications of Skill-Biased Technological Change: International Evidence." *Quarterly Journal of Economics,* November 1998, *113*(4), pp. 1245–79.

Bound, John and Johnson, George. "Changes in the Structure of Wages in the 1980's: An Evaluation of Alternative Explanations." *American Economic Review,* June 1992, *82*(3), pp. 371–92.

Chari, V. V. and Hopenhayn, Hugo. "Vintage Human Capital, Growth, and the Diffusions of New Technology." *Journal of Political Economy,* December 1991, *99*(6), pp. 1142–65.

Ciccone, Antonio. "Falling Real Wages During an Industrial Revolution." Unpublished manuscript, University of California, Berkeley, 1997.

Crafts, Nicholas F. R. *British economic growth during the Industrial Revolution.* Oxford: Clarendon, 1985.

David, Paul A. "Computer and Dynamo: The Modern Productivity Paradox in a Not-Too-Distant Mirror," in *Technology and productivity: The challenge for economic policy.* Paris: OECD, 1991, pp. 315–47.

Davis, Steve J. and Haltiwanger, John. "Wage Dispersion between and within U.S. Manufacturing Plants, 1963–86." *Brookings Papers on Economic Activity: Microeconomics,* 1991, pp. 115–80.

DiNardo, John E. and Pischke, Jörn-Steffen. "The Returns to Computer Use Revisited: Have Pencils Changed the Wage Structure Too?" *Quarterly Journal of Economics,* February 1997, *112*(1), pp. 291–303.

Doms, Mark; Dunne, Timothy and Troske, Kenneth. "Workers, Wages, and Technology." *Quarterly Journal of Economics,* February 1997, *112*(1), pp. 253–90.

Dunne, Timothy; Haltiwanger, John and Troske, Kenneth. "Technology and Jobs: Secular Changes and Cyclical Dynamics." *Carnegie-Rochester Conference Series on Public Policy,* June 1997, *46,* pp. 107–78.

Dunne, Timothy and Schmitz, John A. "Wages, Employment Structure and Employer Size-Wage Premia: Their Relationship to Advanced-Technology Usage at U.S. Manufacturing Plants." *Economica,* February 1995, *62*(245), pp. 89–107.

Galor, Oded and Tsiddon, Daniel. "Technological Progress, Mobility, and Economic Growth." *American Economic Review,* June 1997, *87*(3), pp. 363–82.

Goldin, Claudia and Katz, Lawrence F. "The Origins of Technology-Skill Complementarity." *Quarterly Journal of Economics,* August 1998, *113*(3), pp. 693–732.

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

Griliches, Zvi. "Capital-Skill Complementarity." *Review of Economics and Statistics.* November 1969, *51*(4), pp. 465–68.

Hounshell, David A. *From the American system to mass production: The development of manufacturing technology in the United States.* Baltimore, MD: Johns Hopkins University Press, 1984.

Journal of Economic Perspectives. "Symposium." *Journal of Economic Perspectives,* Spring 1997, *11*(2), pp. 21–96.

Jovanovic, Boyan. "Learning and Growth," in David Kreps and Kenneth Wallis, eds., *Advances in economics.* Cambridge: Cambridge University Press, 1996.

Krueger, Alan B. "How Computers Have Changed the Wage Structure: Evidence from Microdata, 1984–1989." *Quarterly Journal of Economics,* February 1993, *108*(1), pp. 33–60.

Krusell, Per; Ohanian, Lee; Ríos-Rull, José Víctor and Violante, Giovanni L. "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis." Unpublished manuscript, University of Rochester, 1997.

Lindert, Peter H. "Unequal Living Standards," in Roderick Floud and Donald McCloskey, eds., *The economic history of Britain since 1700,* Vol. 1, 2nd Ed. Cambridge: Cambridge University Press, 1994, pp. 357–86.

Lloyd-Ellis, Huw. "Endogenous Technological Change and Wage Inequality." *American Economic Review,* March 1999, *89*(1), pp. 47–77.

Machin, Stephen. "Changes in the Relative Demand for Skills," in Alison Booth and Dennis Snower, eds., *Acquiring skills.* Cambridge: Cambridge University Press, 1996, pp. 127–46.

Mansfield, Edwin. *The production and application of new industrial technology.* New York: Norton, 1977.

Mokyr, Joel. *The lever of riches: Technological creativity and economic progress.* New York: Oxford University Press, 1990.

Mowery, David C. and Rosenberg, Nathan. *Technology and the pursuit of economic growth.* Cambridge: Cambridge University Press, 1989.

Schmitt, John. "The Changing Structure of Male Earnings in Britain, 1974–1988," in Richard B. Freeman and Lawrence F. Katz, eds., *Differences and changes in the wage structure.* Chicago: University of Chicago Press, 1995, pp. 177–204.

U.S. Bureau of the Census, Department of Commerce. *Surveys of manufacturing technology.* Washington, DC: U.S. Government Printing Office, 1988, 1993.

Violante, Giovanni L. "Technological Progress and Skill Dynamics: A Solution to the Wage Dispersion Puzzle?" Unpublished manuscript, University College London, 1997.

von Tunzelmann, Nick. "Atlas of Industrializing Britain 1780–1914: Coal and Steam Power," in John Langton and R. J. Morris, eds., *Atlas of industrializing Britain, 1780–1914.* London: Methuen, 1986, pp. 72–79.