Growth Accounting with Technological Revolutions

Andreas Hornstein

G enerally, technological progress proceeds at a slow and measured pace, with only incremental improvements seen in existing products and technologies in the economy. At times, however, the pace accelerates, and the economy experiences a technological revolution during which radically new products and technologies are introduced. Recent discussions suggest that the world economy is currently experiencing just such a revolution, or paradigm shift, and that this revolution accounts for some of the observed decline and rebound of productivity growth. For example, David (1991) argues that the effect of information technologies on today's economy is comparable to the effects of the introduction of the dynamo and the subsequent availability of electric power in the late-nineteenth and early-twentieth centuries. It is important to understand the effects of technological progress as reflected in productivity growth because productivity growth determines the economy's long-run growth of output, consumption, and factor income such as wages.

In this article I consider one particular parable of a paradigm shift. This story builds on three assumptions: first, that technological change is associated with the introduction of new goods, in particular that new technologies are embodied in new machines; second, that production units learn about the newly introduced technologies, that is, new technologies do not immediately attain their full productivity potential, but instead productivity increases gradually for some time; and third, that the experience which production units have with existing technologies affects their ability to adopt new technologies.[1]

[1] The ideas expressed in this article are based on work by Greenwood, Hercowitz, and Krusell (1997), Hornstein and Krusell (1996), Greenwood and Jovanovic (1998), and Greenwood and Yorukoglu (1997).

In the following pages I summarize the available evidence in support of these assumptions and then speculate on the possible implications of a paradigm shift for future output and productivity growth based on a parametric version of the standard neoclassical growth model. I find that all three assumptions together can account for a substantial and long-lasting decline in measured productivity and output growth during the initial stages of a technological revolution. This initial period is then followed by a long period of above-average long-run growth. Unfortunately, the results depend crucially on how experience with existing technologies affects the ability to adopt new technologies, a feature of the economy about which we know very little. An alternative parameterization of this feature of the economy predicts that the effects of a technological revolution on productivity and output growth might be negligible. Finally, I reconsider the evidence on the slowdown of measured productivity growth and find that it appears to be less dramatic if we calculate real output numbers using a more reliable price index.

1.   SOME EVIDENCE ON TECHNOLOGICAL CHANGE

The Rate of Capital-Embodied Technological Change has Accelerated in the Early '70s

When people talk about a new technological revolution, they usually refer to the more widespread use of computers: the application of computers makes new products and services possible, it changes the way production processes are organized, and it is no longer limited to a small fraction of the economy. Unfortunately, many of these observations are anecdotal and provide only limited quantitative support for the impact of computers on the economy. There is one observation, however, that we all make and that might well be quantified; namely, that each new generation of PCs tends to do more things faster than the previous generation, yet we do not have to pay more for these higher-quality PCs. In short, for PCs the price-per-quality unit has been declining at a dramatic rate. This observation applies not only to PCs but to many other products, particularly producer-durable goods such as new capital goods.

While it is easy to say that new products are of better quality, it is difficult to actually measure and compare quality across different goods. In an extensive study, Gordon (1990) has constructed measures of the price of producer-durable equipment that account for quality changes. The line labeled $1/q$ in Figure 1 graphs the price of new producer-durable equipment relative to the price of nondurable consumption for the postwar U.S. economy.[2] I identify the rate of

---

[2] The series on the relative price of producer-durable equipment is from Greenwood, Hercowitz, and Krusell (1997). I have extrapolated the series from 1990 on using information on the price of producer-durable equipment from the National Income Accounts. Consumption covers nondurable goods and services, excluding housing services. Hornstein (1999) provides a complete description of the data used.

Figure 1   Measures of Embodied and Disembodied Technological Change
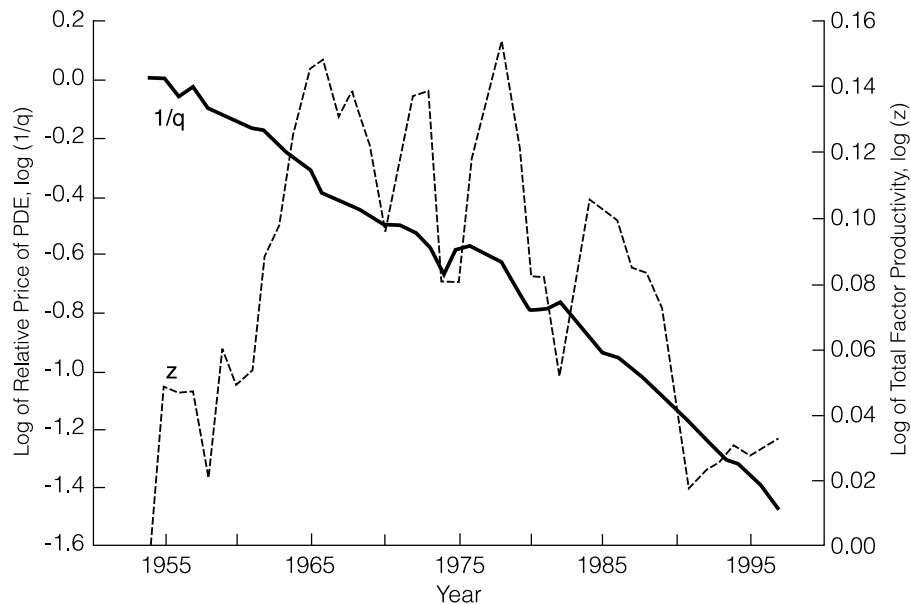
price decline with increased productivity in the capital goods producing sector that is embodied in the new capital goods.[3] In this figure it is apparent that producer-durable equipment goods have become cheaper over time relative to consumption goods and that the rate of price decline has accelerated in the mid-'70s from 3 percent before 1973 to 4.3 percent after 1977. A substantial part of the accelerated rate of price decline can be attributed to the fact that information technologies have gained more widespread application in the design of producer-durable equipment.

Learning-by-Doing is an Important Feature of Production

New products or new plants do not attain their full potential at the time they are introduced. Rather, we find that for some period of time productivity for

---

[3] In general, relative prices may change because the technology changes (shift of the production possibility frontier, PPF) or because of simple substitution between goods (movements along a PPF). Notice, however, that with an unchanged technology we would expect the relative price of a good to fall only if relatively less of the good is produced. Yet we have not observed a decline in the investment rate that should correspond to the decline in the relative price of capital. Work that tries to account for substitution effects finds even more acceleration in the rate of capital-embodied technological change (Hornstein and Krusell 1996).

a new good or plant is increasing. This increase in productivity is attributed to *learning-by-doing* (LBD); that is, firms acquire experience and improve their efficiency in resource use in the process of producing a good. One can think of this process as the accumulation of informational capital. This LBD phenomenon is so widespread and uniform across industries that the management literature summarizes it with the "20 percent rule," according to which labor productivity increases by 20 percent for every doubling of cumulative production (see, e.g., Hall and Howell [1985]).

One of the most frequently cited LBD examples is the case of the liberty ships of World War II. The more ships a navy yard built, the smaller was the labor input required for the next vessel it built (Figure 2). A more recent example of LBD is the production of dynamic random access memory (DRAM) chips in the semiconductor industry. Figure 3 displays the time paths for the average unit price and total shipments of successive generations of DRAM chips. This figure displays two common features of LBD. First, productivity improvements during the early stages of production are dramatic. Second, these improvements are attained within a short period of time, occurring within the first three to five years of production. Indeed, most of the productivity improvements have been made once shipments of a chip generation reach their peak. Notice also that during the first few years a new generation of chips is produced, the unit price is higher than the one of the previous generation.[4] The DRAM chip example also points to an important feature for my discussion of accelerated capital-embodied technological change: How much of the experience accumulated in the production of one generation of DRAM chips can be transferred to the production of the next generation of chips? More generally, how much of the experience accumulated for existing technologies can be applied to new technologies? The answer to this question is still open. Evidence from the semiconductor industry indicates that the transfer of experience is limited (Irwin and Klenow 1994).

New Technologies Diffuse Slowly Through the Economy

When a radically new technology becomes available, not everybody in the economy will adopt this new technology simultaneously. For some time the use of the old and new technology will coexist while firms continue to make improvements in the old technology. This situation will occur since there are costs to adopting new technologies such as learning costs. Potentially, a new technology may be much more productive than the old technology, but initially users of the new technology have to start with a low experience level relative to that of old technologies.

---

[4] My interpretation that a decline in average unit price reflects an increase in productivity should be taken with a grain of salt since the market structure in the semiconductor industry is only approximately competitive.

Figure 2   Reductions in Man-Hours per Vessel with Increasing Production
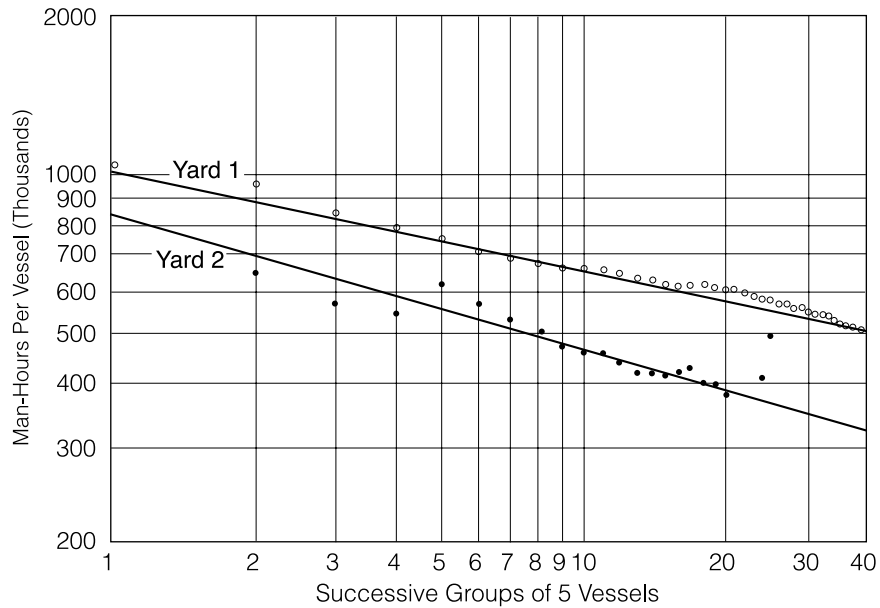

Source: Lucas (1993).

The idea that new technologies diffuse slowly through the economy relates to the observation that the use of new products diffuses slowly through the economy. Experts have made this observation for a wide variety of products from diesel locomotives to DRAM chips (Figures 3 and 4).[5] David (1991) makes a similar observation on the diffusion of the use of electrical power in the late-nineteenth and early-twentieth centuries.

2.   TECHNOLOGY DIFFUSION IN A SIMPLE MODEL OF CAPITAL-EMBODIED TECHNOLOGICAL CHANGE AND LEARNING

The appearance of a new technology in the economy can significantly affect output and productivity growth during the transitional period when the new technology replaces the old technology. These effects come about because learning introduces another kind of capital that is not measured, *informational*

---

[5] Figure 3b plots shipments of DRAM chips from different generations, and Figure 4 plots the numbers of diesels in use as a fraction of the total number of locomotives.

Figure 3   Dynamic Random Access Memory Semiconductors

a. Prices

b. Shipments
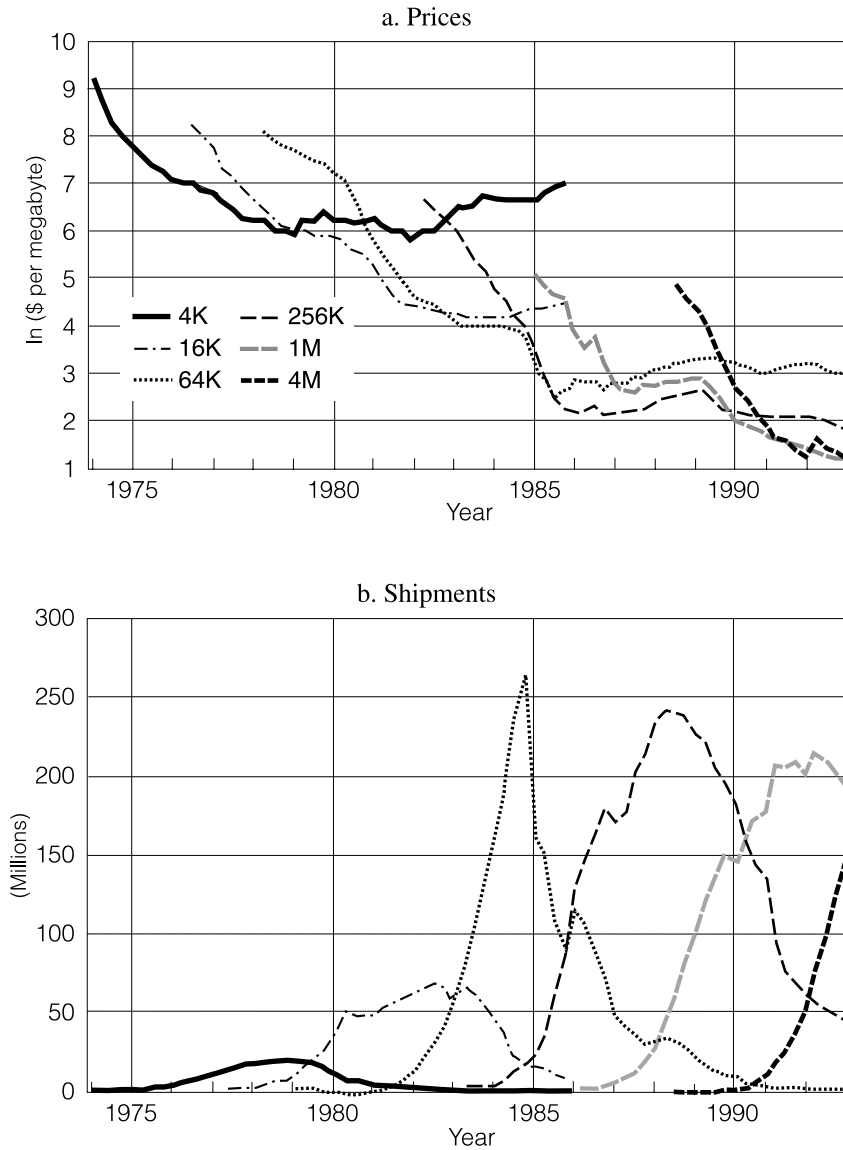
capital, and during the transitional period this capital stock can change signifi-
cantly. This change in informational capital has real output growth effects, and
it creates a problem for measuring productivity growth.

**Figure 4   Diesel Locomotion in the U.S. Railroad Industry,
1925–66: Diffusion**
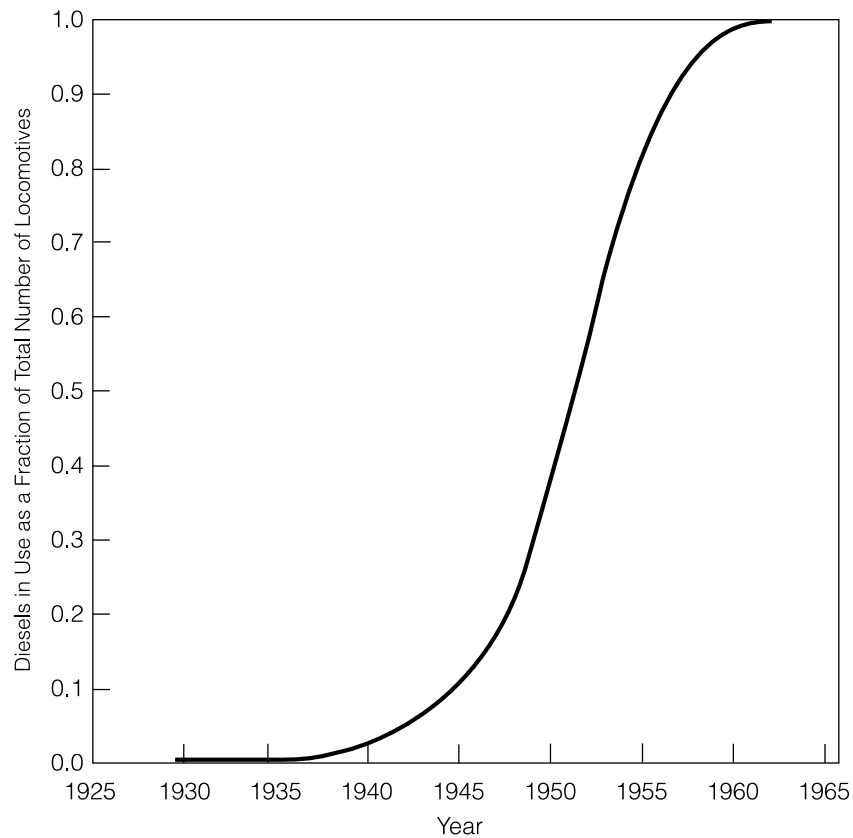
Source: Jovanovic and McDonald (1994).

Informational capital represents the economy's experience with various vintages of capital goods, and it is not part of our standard measure of capital. Consequently, we do not measure changes of informational capital that occur during transitional periods. In particular, after we correct for depreciation, we assign the same value to capital from different vintages. So during transitional periods when substantial investment in new technologies with lower experience occurs, we tend to overestimate the contribution to output from investment in these new technologies. Because we overestimate capital accumulation, we underestimate total factor productivity growth. There is also a real effect of learning, since output growth slows down in the transitional period. A feature of this learning is that during the transitional period, production with new technologies is relatively less efficient than production with old technologies.

In the next section I will try to quantify the implications for output and productivity growth measurement when a new technology is introduced in a simple vintage capital model with learning. The structure of the model is very mechanical and many of the elements discussed above are taken as exogenous.

The Solow Growth Model and Growth Accounting

I will start with the standard Solow growth model, which assumes a neoclassical production structure and a constant savings and investment rate. Each period, a homogeneous good y_t is produced using a constant-returns-to-scale technology with inputs capital k_t and labor n_t,

$$y_t = z_t k_t^\alpha n_t^{1-\alpha}, \tag{1}$$

where the elasticity of output with respect to capital satisfies $0 < \alpha < 1$. For simplicity I have assumed a Cobb-Douglas production function. Technological change is represented through changes in total factor productivity z_t and is disembodied; that is, with the same inputs, output increases when total factor productivity (TFP) increases. The economy's endowment of labor is fixed, $n_t = 1$. The output good can be used for consumption c_t or investment i_t:

$$c_t + i_t = y_t. \tag{2}$$

Investment is used to augment the capital stock and capital depreciates at a constant rate δ:

$$k_{t+1} = (1 - \delta)k_t + i_t, \tag{3}$$

and $0 < \delta < 1$. Expenditures on investment are assumed to be a constant fraction σ of output,

$$i_t = \sigma y_t, \tag{4}$$

and $0 < \sigma < 1$.

Assume that TFP grows at a constant rate, $z_{t+1} = \gamma_z z_t$ and $\gamma_z \geq 1$. It can be easily verified that an equilibrium exists for this economy where output, consumption, investment, and the capital stock all grow at constant rates. Such an equilibrium is called a *balanced growth path*. For the following let g_x denote the gross growth rate of the variable x: that is, $g_x = x_t/x_{t-1}$. From the savings equation (4), it follows that if both investment and output grow at a constant rate, then they must grow at the same rate, $g_y = g_i = g$. In turn, the resource constraint (2) shows that consumption must grow at that same rate $g_c = g$. Dividing the capital accumulation equation (3) by the capital stock k_t subsequently shows that if the capital stock grows at a constant rate, it must grow at the same rate as investment, $g_k = g$. Finally, the production function (1) relates the economy's output growth rate to the growth of inputs and the exogenous

productivity growth rate $g = g_y = \gamma_z g_k^{\alpha} = \gamma_z g^{\alpha}$.[6] From this expression one can see that the economy's growth rate on the balanced growth path increases with the productivity growth rate and with the capital elasticity of output,

$$g = \gamma_z^{1/(1-\alpha)}. \tag{5}$$

We know that TFP in this economy is z_t, but how can we measure TFP if we do not observe z_t? In order to calculate the percentage change of TFP, take the log of equation (1), take the first difference,[7] and solve for the TFP growth rate $\hat{z}$,

$$\hat{z}_t = \hat{y}_t - \alpha \hat{k}_t - (1 - \alpha)\hat{n}_t.$$

Here the measure of TFP growth requires observations on the growth rates of output and inputs and knowledge of the parameter α. Solow's (1957) important insight was that in a competitive economy α can be measured through observations on factor income shares. Suppose that all markets in this economy are competitive and that everybody has access to the technology represented by (1). Then consider a firm that maximizes profits, sells the output good at a price p_t, and hires labor (capital) services at the wage rate w_t (capital rental rate u_t). In order to maximize profits, the firm will hire labor (capital) services until the marginal revenue from the last unit of labor (capital) services hired equals its price:

$$p_t MPN_t = p_t (1 - \alpha) z_t k_t^{\alpha} n_t^{-\alpha} = w_t \tag{6a}$$

$$p_t MPK_t = p_t \alpha z_t k_t^{\alpha-1} n_t^{1-\alpha} = u_t. \tag{6b}$$

Multiplying each side of the equation with $n_t/p_t y_t$ ($k_t/p_t y_t$) shows that the labor (capital) coefficient in the production function equals the share in total revenues that goes to labor (capital):[8]

$$1 - \alpha = w_t n_t/p_t y_t = s_{nt}$$

$$\alpha = u_t k_t/p_t y_t = s_{kt}.$$

---

[6] From equations (1), (3), and (4), it follows that a balanced growth path is associated with a particular level of the capital stock in the initial period k_0. One can show that the economy converges toward this balanced growth path if it starts with a different level of capital.

[7] In the following, a hat denotes the net growth rate of a variable: for example, $\hat{x}_t = (x_t - x_{t-1})/x_t$. For small changes in a variable, the first difference of the logs approximates the growth rate; for example, $\hat{x}_t = \ln x_t - \ln x_{t-1}$.

[8] Since the two coefficients sum to one, total payments to the two production factors capital and labor exhaust revenues; that is, there are zero profits. This is not specific to the assumption of a Cobb-Douglas production function. In general, profits are zero when production is constant returns to scale and all markets are competitive.

We can therefore measure productivity growth using observations on output growth, input growth, and factor income shares. This measure of TFP growth is the *Solow residual*:

$$\hat{z}_t^m = \hat{y}_t - s_{kt}\hat{k}_t - s_{nt}\hat{n}_t = \hat{z}_t. \tag{7}$$

The Solow residual provides an accurate measure of disembodied technological change not only for a Cobb-Douglas production structure but for any constant-returns-to-scale economy, as long as we are willing to assume that all markets are competitive. Finally, note that the wage and capital rental rate equations (6a and 6b) also imply that on a balanced growth path real wages w_t/p_t will grow at the economywide growth rate g, which is determined by the productivity growth rate, and that the real rental rate of capital is constant.

Capital-Embodied Technological Change

The secular decline of the relative price of producer-durable goods suggests that a substantial part of technological progress is embodied in new capital goods. A straightforward modification allows me to account for capital-embodied technological change in the Solow growth model. In the model described above the homogeneous output good can be used for consumption or investment, and the marginal rate of transformation between consumption and investment goods is fixed. In particular I have assumed that one unit of the consumption good can be transformed into q_t units of the investment good and $q_t = 1$. In order to show that over time the economy becomes more efficient in the production of capital goods, I simply assume that over time q_t grows at a constant rate, $q_{t+1} = \gamma_q q_t$ and $\gamma_q \geq 1$. The resource constraint for the output good is now

$$c_t + i_t/q_t = y_t. \tag{2a}$$

At the same time that the economy becomes more efficient in the production of capital goods, the relative price of capital goods $1/q_t$ declines. I continue to measure output in terms of consumption goods and assume that expenditures on investment goods in terms of consumption goods represent a constant fraction of income:

$$i_t/q_t = \sigma y_t. \tag{4a}$$

Analogous to the previous economy, there is a balanced growth path where output, consumption, investment, and capital all grow at constant rates:

$$g_y = g_c = (\gamma_z\gamma_q^\alpha)^{1/(1-\alpha)} \text{ and } g_i = g_k = (\gamma_z\gamma_q)^{1/(1-\alpha)}. \tag{5a}$$

The measurement of TFP, that is, disembodied technological change, is affected in two ways by the presence of capital-embodied technological change. First, the capital stock measure is constructed as the cumulative sum of unde-preciated past investment based on equation (3). Since changes in the quality of

new capital goods are the hallmark of embodied technological change, we have to use an appropriate price index that accounts for these quality changes when we deflate nominal investment series to obtain real investment expenditures. Second, because the relative price between consumption and investment goods is changing over time, we have to decide whether we want to measure output in terms of consumption or investment goods. Since ultimate well-being in the economy depends on the availability of consumption goods, I decide to measure output in terms of consumption goods. The line labeled z in Figure 1 displays the measured TFP levels for the postwar U.S. economy. Here the measured capital stock is adjusted for embodied technological change using data on the relative price of durable goods.[9] Notice that contrary to capital-embodied technological change, which was positive for all of the postwar period, measured TFP does not represent a success story for the U.S. economy. Although TFP was increasing rapidly in the late '50s and '60s, TFP stagnated in the early '70s and has actually declined since the mid-'70s when the rate of embodied technological change accelerated. Recently, starting in the '90s, there has been a slight recovery of TFP, but the apparent negative trend in the '70s and '80s seems hard to rationalize.

Learning and Growth Accounting

The observed decline in measured TFP could simply be due to measurement error; that is, there never was a decline in actual TFP. To make sense of this explanation I provide a candidate for what has been mismeasured, and I argue why the measurement problem got worse in the mid-'70s and why we now observe a trend reversal. I suggest that the effective stock of capital has been mismeasured. In particular, I consider the possibility that standard measures of capital do not include informational capital in the economy. In the following I introduce informational capital into the Solow growth model through a simple model of learning. I show that even though measured capital does not include informational capital, there is no measurement problem on the balanced growth path; the measured capital stock may overestimate the effective capital stock during transitional periods when there are significant changes in the economy's informational capital stock.

Assume that new capital goods do not immediately attain their full potential, but in the process of producing goods, more is learned about each capital good and the efficiency with which it is used increases over time. We now have to distinguish between different vintages of capital goods because a producer has less experience with a capital good that is newly introduced than with a capital good that has been around for some time. Let $k_{t,a}$ denote a capital good

---

[9] The measure of TFP is based on work by Greenwood, Hercowitz, and Krusell (1997) as extended in Hornstein (1999). For a more detailed description see either of the two references.

that is a years old at time t. If this capital good is employed with $n_{t,a}$ units of labor, output $y_{t,a}$ is

$$y_{t,a} = z_t e_{t,a} k_{t,a}^{\alpha} n_{t,a}^{1-\alpha}, \tag{1a}$$

where $e_{t,a}$ is the experience index of a capital good that is a years old. For simplicity I assume that maximal experience is one and convergence to it is geometric at rate λ:

$$1 - e_{t+1,a+1} = \lambda(1 - e_{t,a}) \text{ for } a = 1, 2, \ldots, \tag{8}$$

starting from some initial experience level $0 \le e_{t,1} \le 1$, and $0 < \lambda < 1$. I continue to assume that capital depreciates at rate δ:

$$k_{t+1,a+1} = (1 - \delta)k_{t,a}. \tag{3a}$$

Total output, employment, and investment are

$$y_t = \sum_{a=1}^{\infty} y_{t,a}, n_t = \sum_{a=1}^{\infty} n_{t,a}, \text{ and } i_t = k_{t+1,1}, \tag{9}$$

and I continue to assume that the markets for output, labor, and the different capital vintages are all competitive. An attractive feature of this model is the existence of an exact aggregate capital index. We can write aggregate output as a Cobb-Douglas function of total employment and the aggregate capital index $\bar{k}$:[10]

$$y_t = z_t \bar{k}_t^{\alpha} n_t^{1-\alpha}, \text{ and } \bar{k}_t = \sum_{a=1}^{\infty} e_{t,a}^{1/\alpha} k_{t,a}. \tag{9a}$$

From this expression one can see how informational capital, $e_t = \{e_{t,a} : a = 1, 2, \ldots\}$, affects aggregate output. Note that the usual measure of the

---

[10] The aggregate capital index can be derived as follows. A profit-maximizing competitive firm using vintage a capital goods hires labor until it equates the marginal revenue of labor with its marginal cost, $p(1 - \alpha)ze_a k_a^{\alpha} n_a^{-\alpha} = w$. Solving this expression for n_a defines the demand for labor by firms using vintage a capital, $n_a = [(1 - \alpha)ze_a/(w/p)]^{1/\alpha} k_a$. The real wage w/p then adjusts such that the total demand for labor is equal to the supply of labor:

$$n = \sum_{a=1}^{\infty} n_a = [(1 - \alpha)z/(w/p)]^{1/\alpha} \sum_{a=1}^{\infty} e_a^{1/\alpha} k_a = [(1 - \alpha)z/(w/p)]^{1/\alpha} \bar{k}.$$

One can solve this expression for the equilibrium real wage, substitute it in the labor demand equation, and obtain the output of firms using vintage a capital as $y_a = ze_a^{1/\alpha} k_a (n/\bar{k})^{1-\alpha}$. Total output is then

$$y = \sum_{a=0}^{\infty} y_a = z\bar{k}^{\alpha} n^{1-\alpha}.$$

economy's capital stock as the sum of undepreciated past investment does not take into account the informational capital

$$k_t^m = \sum_{a=1}^{\infty} k_{t,a} = (1 - \delta)k_{t-1}^m + i_{t-1}. \tag{9b}$$

To close the model I identify what determines initial experience with a new capital good. I assume that there is an externality, and experience with older capital goods is partially transferrable to new capital goods according to the following expression:

$$e_{t+1,1} = \frac{\theta}{1 - \rho} \sum_{a=1}^{\infty} \rho^{a-1} n_{t,a} e_{t,a}, \tag{8a}$$

with $0 < \rho < 1$ and $\theta > 0$. This formulation of the learning externality follows Lucas (1993). The factor ρ^a measures the extent to which experience with vintage a contributes to initial experience with new capital goods. The larger that ρ is, the more important is experience with existing capital goods. Since $\rho < 1$, experience with older vintages is less important for the initial experience with a new capital good. Notice also that I have assumed the contribution of vintage a is weighted by how intensively this vintage is used, whereby I measure the intensity of use by the share in employment.

The balanced growth path of this economy is very similar to the path of the previous economy. Output, consumption, investment, and capital grow at the same rates, and the initial experience e_1 is constant. Because the initial experience is constant, the informational capital does not change, $e_{ta} = e_a$, and the exact aggregate capital index (9a) and the measured capital stock (9b) grow at the same rate. Therefore, the Solow residual accurately reflects true growth of TFP. If the economy is not on the balanced growth path, three things happen. Initial experience and the informational capital changes over time, changes in the measured capital stock do not accurately reflect changes in the exact aggregate capital index, and the Solow residual mismeasures true TFP growth.

The economy may not be on its balanced growth path for various reasons. Here I consider the possibility that the acceleration of capital-embodied technological change in the mid-'70s was associated with a qualitative change in the kind of technology used. Furthermore, the adoption of this new technology proceeded gradually. To be more specific assume that at some time t_0 this new qualitatively different technology becomes available. From this point on I distinguish between vintages belonging to the old technology, $i = 1$, and vintages belonging to the new technology, $i = 2$. This means that in any time period t output, capital, employment, and experience are now indexed by the type of technology i and its vintage a, $\{y_{t,a}^i, k_{t,a}^i, e_{t,a}^i, n_{t,a}^i\}$. I assume that the new technology is potentially better because capital-embodied technological progress

proceeds at a higher rate for the new technology $\gamma_q^2 > \gamma_q^1$ and $q_{t_0}^2 = q_{t_0}^1$. At first, however, the new technology may be worse because the economy has less experience with it. Since the new technology may be initially inferior, I assume that the new technology diffuses slowly. In particular, only a fraction ψ_t of total investment expenditures is used for the purchase of capital goods with the new technology, and

$$
\psi_t \begin{cases} = 0 & \text{for } t < t_0, \\ \in (0,1) & \text{for } t = t_0 + 1, \ldots, \ t_0 + T, \\ = 1 & \text{for } t > t_0 + T, \end{cases} \tag{10}
$$

and ψ_t increases monotonically. As before, initial experience e^i for a new vintage of a technology i depends on the existing experience with older vintages of that technology,

$$
e_{t+1,1}^i = \frac{\theta}{1 - \rho} \sum_{a=1}^{\infty} \rho^{a-1} n_{t,a} e_{t,a}^i. \tag{8b}
$$

For completeness assume that the experience of a new technology vintage that never existed is zero; that is, $e_{ta}^2 = 0$ for $t - a < t_0$.[11]

   In order to consider the quantitative implications of the diffusion of a new technology, I select parameter values for the economy that are consistent with observations on long-run growth, the evidence on the accelerated embodied technological change, learning in the economy, and the diffusion of new technologies.

   In the postwar U.S. economy, the average annual depreciation rate is about 10 percent, the average investment rate is about 20 percent, and the average capital income share is about 30 percent. I assume that there is no disembodied technological change such that we can interpret the output and measured TFP growth rates as possible losses/gains due to the diffusion of a new technology. I also assume that the new technology is implemented beginning in 1974 and that it will take 40 years for all new investment to take the form of the new technology. This means that we have passed the midpoint of the diffusion process. The parameterization of the diffusion process (T, ψ_t) is consistent with observations as discussed in Section 2. The rate of capital-embodied technological change for the old and new technology corresponds to the average rate of decline for the relative price of equipment before and after 1974. The parameterization of the internal learning process $(\lambda, e_{t_0,1}^1)$ is based on Bahk and Gort (1993). We know the least about learning externalities (ρ, θ). I simply assume that $\rho = 0.8$ and that in the years before 1974 the economy was on its balanced growth path. With this observation I can recover the value of θ.

---

[11] The assumption that experience is not transferable across technologies is extreme, but allowing for partial transferability changes the results insignificantly.
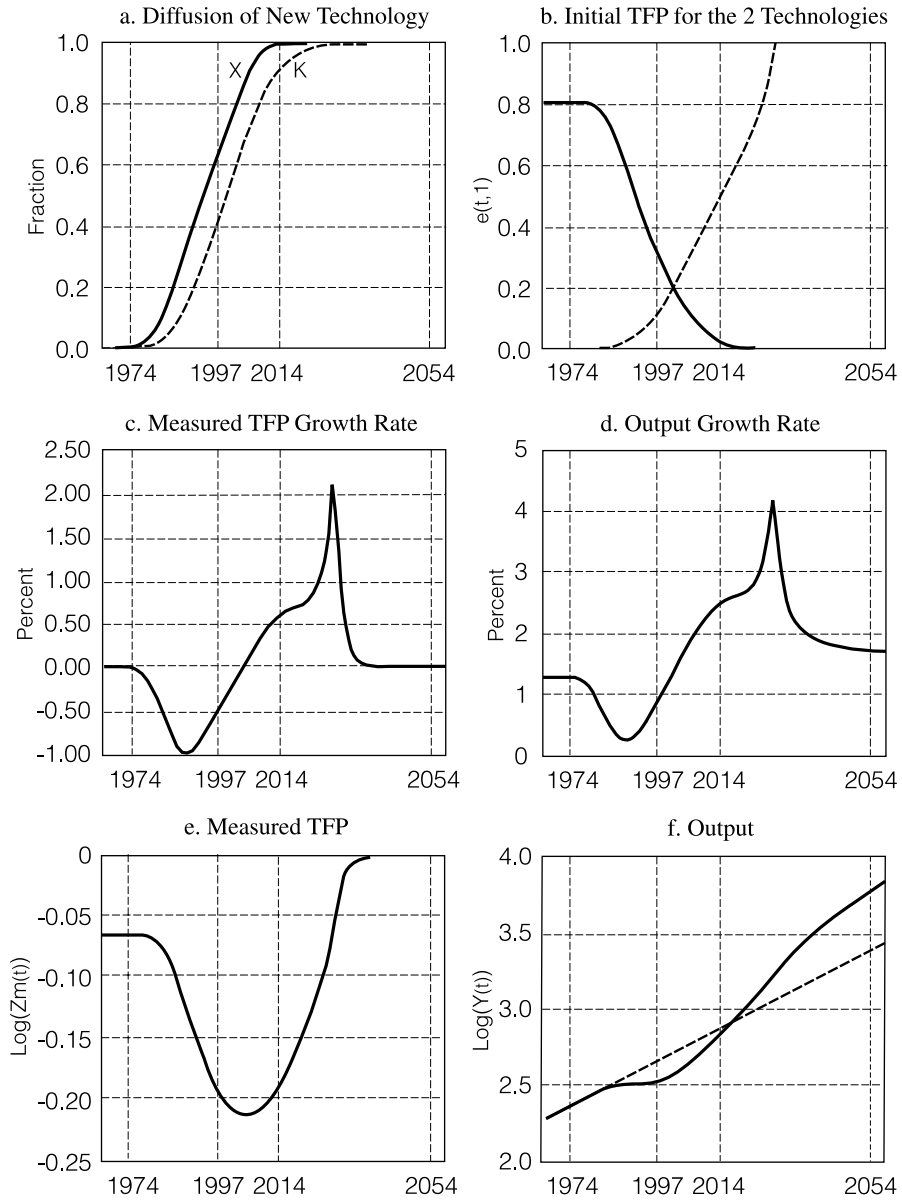
Table 1 Parameter Values

Solow growth model	$\alpha = 0.3$, $\delta = 0.1$, $\sigma = 0.20$
Disembodied technological change	$\gamma_z = 1.00$
Capital-embodied technological change	$\gamma_q^1 = 1.03$ and $\gamma_q^2 = 1.04$
Learning	$\lambda = 0.7$ and $e_{t_0,1}^1 = 0.8$
Learning externality	$\rho = 0.8$ and $\theta = 12.11$
Diffusion	$T = 40$ and ψ_t follows an S-shaped diffusion (third-order polynomial)

The results are displayed in Figure 5. Panel a shows the gradual diffusion of the new technology for investment and the capital stock. Since investment adds to the existing capital stock, the diffusion of the new technology in the total capital stock proceeds at a slower rate than it does relative to investment. Panels c and d display measured TFP growth rates and output growth rates, and we observe a long-lasting and substantial decline in measured TFP growth and output growth (1 percentage point). This decline bottoms out in the mid-'80s, and we are now in a recovery phase. According to this simulation we can expect a considerable increase of the trend growth rates for measured TFP and output for the next 20 years. Panel f shows that the effects of the lower output growth are quite substantial in the sense that another 15 years have to pass before the level of output catches up with the initial balanced growth path.[12]

Why do we get these big effects during the transitional period when the new technology is adopted? The simple answer lies in the graph of initial experience for the two technologies (panel b of Figure 5). Notice that initial experience in the old technology is declining during the transitional phase. The decline occurs because according to the specification of the learning externality (8a), the contribution of a vintage is weighted by employment in that vintage. During the transitional phase employment shifts from old to new technologies and, with this learning specification, the economy tends to "forget" about the old technology. Sizeable changes in output and measured TFP growth do not occur, however, because investment in new technologies starts out with a low experience; these changes make up only a small fraction of total investment after all. Rather the big changes in output and measured TFP growth occur because initial experience is falling for investment in old technologies, and this investment contributes the most to total capital accumulation.

---

[12] The results are sensitive with respect to technology spillovers ρ. If spillovers are unimportant ($\rho = 0.5$), then the decline in measured TFP growth is much more persistent, and output growth does not overshoot very much.

Figure 5   A Transition Path with Big Effects

a. Diffusion of New Technology

b. Initial TFP for the 2 Technologies

c. Measured TFP Growth Rate

d. Output Growth Rate

e. Measured TFP

f. Output

We can evaluate the importance of this effect by changing the specification of the learning externality (8a) such that we do not weight experience by employment; that is, how much the experience of an old vintage contributes to the initial experience of a new vintage is independent of how intensively the old vintage is used:

$$e_{t+1,1}^i = \frac{\theta}{1-\rho} \sum_{a=1}^{\infty} \rho^{a-1} e_{t,a}^i. \tag{8c}$$
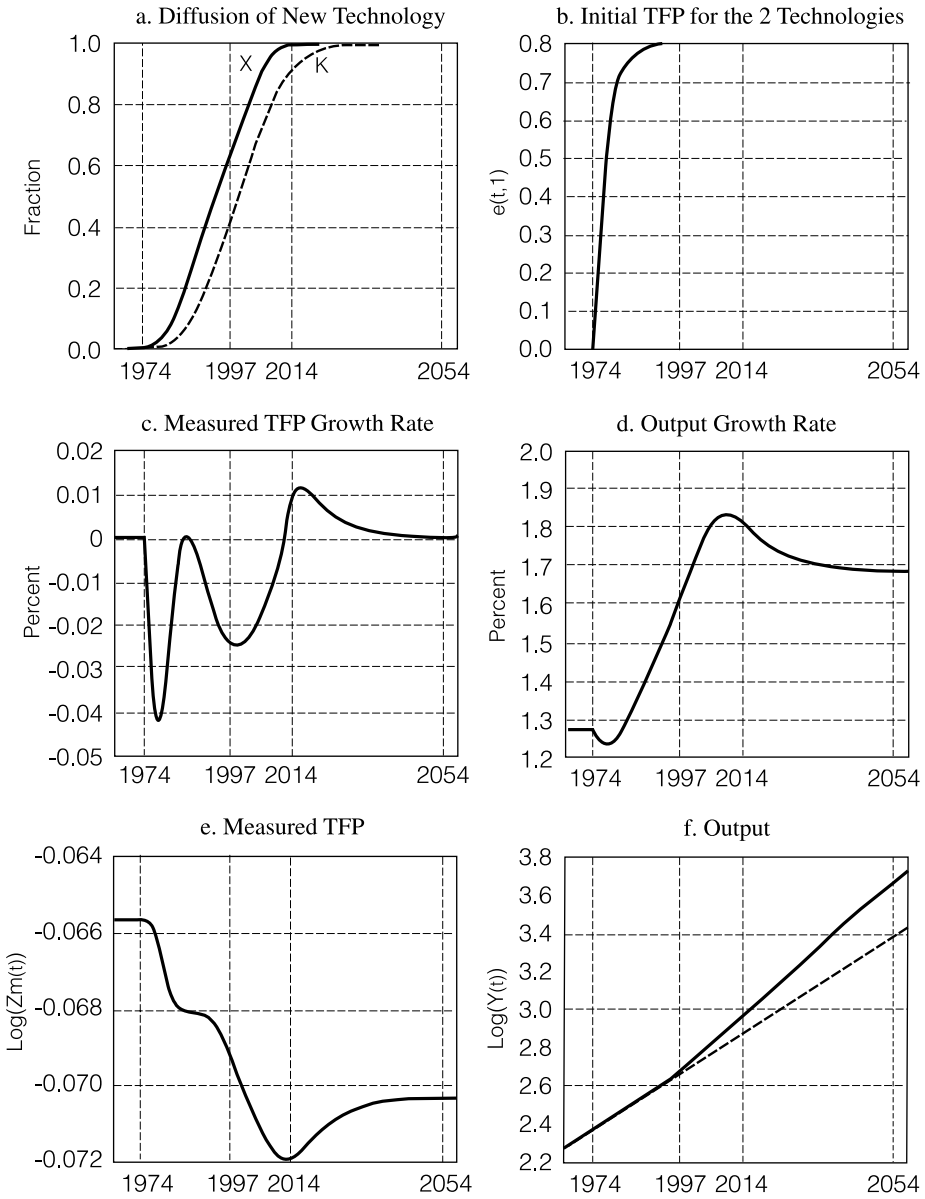
The results of this change are displayed in Figure 6. Note that with this specification initial experience with the old technology remains constant at 0.8. As we can see, the maximal reduction in measured TFP growth is now only 0.04 percentage points, as opposed to 1 percentage point previously, and there is almost no decline in output growth; the maximal increase corresponds to the balanced growth increase of about 0.5 percentage points.

I am not aware of any empirical work that has studied the quantitative properties of the transfer of knowledge in the economy and that would allow us to pick between the two learning specifications (8a) and (8c). Although I find specification (8a) reasonable—in the sense that intensity of use should matter for the transfer of knowledge—and although it is quite possible that an economy "forgets" about old technologies if they are not used, I do not believe that the process occurs as fast as implied by the specification above. If, as I believe, the economy is not that forgetful, then specification (8c) may be a good short- to medium-term approximation, and I would have to conclude that the possible effects of a technological revolution are limited.

3. RECONSIDERING THE MEASUREMENT OF PRODUCTIVITY GROWTH

This article reviews the possible implications of a technological revolution for the measurement of the U.S. economy's productivity performance. I have shown evidence for the acceleration of capital-embodied technological change and at the same time a substantial decline of TFP, which represents disembodied technological change. I have argued that part of the decline in TFP can in principle be attributed to a measurement problem associated with accumulating informational capital during a technological revolution. Unfortunately, the process by which informational capital is accumulated in an economy is not well understood, and any exercise that studies this aspect of the economy has to be somewhat speculative in nature. I would like to conclude my discussion of the U.S. economy's productivity performance with one more observation. Although this observation makes the description of productivity behavior even more ambiguous, it seems to indicate that the performance of the U.S. economy has not been as bad as Figure 1 suggests.

Figure 6   A Transition Path with Small Effects

a. Diffusion of New Technology

b. Initial TFP for the 2 Technologies

c. Measured TFP Growth Rate

d. Output Growth Rate

e. Measured TFP

f. Output

My discussion of the implications of a technological revolution has focused on problems associated with the measurement of capital in a broad sense. Part of the measurement problem is accounting for changes in the quality of producer-durable goods, but for this part I have taken the view that Gordon's (1990) price index does account for most of the quality changes that occur for producer-durable goods. I have also identified embodied technological progress with the rate of decline of the price of producer-durable goods relative to consumption goods. At this point I should note that the quality of consumption goods also changes over time, a process that in principle is no different from that of producer-durable goods. But this means that for the construction of a consumer price index one also has to be careful how one accounts for quality change in new consumer goods. To the extent that our consumer price index does not capture quality changes in goods, we will overestimate the rate of price increase and underestimate the growth in real consumption.[13]

The diffusion of information technologies has certainly affected the quality of consumer goods we are now able to purchase, an observation that is most evident for consumer services. Take, for example, the services provided by the financial sector: we are now able to obtain cash at conveniently located automatic teller machines, we can access our bank accounts and make transactions from home, we can trade shares directly on the Internet without going through a broker, etc. It has always been recognized that accounting for quality changes is relatively more difficult for services than it is for commodities, a problem that has probably been exacerbated through increasingly widespread use of the new information technologies.[14]

A price index that overestimates the rate of price increase for consumer goods has two implications for the productivity growth measures I have discussed in this article. First, since the rate of decline for the price of producer-durable goods relative to the price of consumer goods is overestimated, the rate of embodied technological change is overestimated. Second, because output as measured in terms of consumption goods is actually growing faster than the consumption price index seems to indicate, the rate of disembodied technological change is underestimated. Can we say anything about the potential magnitude of this measurement problem?

I have argued that the measurement problem is probably more relevant for the consumption of services rather than the consumption of goods. If services made up only a small fraction of consumption, the potential bias would probably be small, but today expenditures on services excluding housing are about 50 percent higher than expenditures on nondurable goods. Since the price index for nondurable consumption goods appears to be less subject to measurement

---

[13] For a discussion of the potential biases in the consumer price index, see Boskin et al. (1996).

[14] See Griliches (1994) on the quality of output and price indexes for different industries.

Figure 7   Measures of Embodied and Disembodied Technological Change Reconsidered
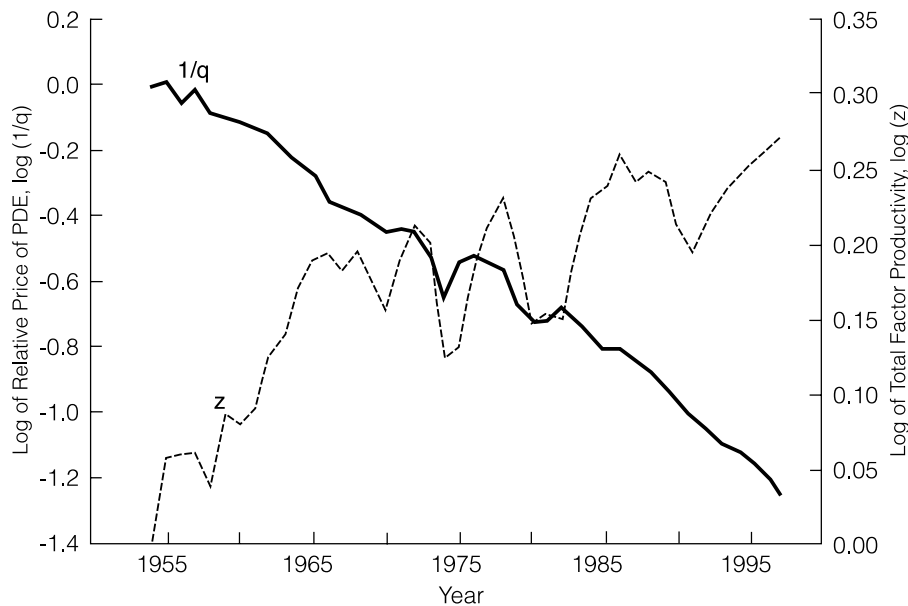
error than the price index for services, I recalculated the estimates for embodied and disembodied technological change using the price index for nondurable consumption goods only, rather than the price index for nondurable goods and services (excluding housing) as shown in Figure 1. The revised productivity series are graphed in Figure 7.

The alternative measure of real output mainly effects the measure of disembodied technological change as opposed to the measure of embodied technological change. Embodied technological change now proceeds at a slower rate, and it does not accelerate as much in the mid-'70s.[15] The effect on the measure of disembodied technological change as reflected in TFP growth is more dramatic. With the new measure of real output, TFP growth still stagnates starting in the '70s, but there is no longer a secular decline. Notice also the strong recovery of TFP since the early '90s, although it remains to be seen whether this is a purely cyclical upswing or whether it represents a change in the long-run growth path for TFP. In conclusion, as is evident from Figures 1 and 7, the productivity

---

[15] The rate of price decline now accelerates from 2.7 percent before 1973 to 3.5 percent after 1977.

performance of the U.S. economy appears to be consistent with a wide range of views, from pessimistic to guardedly optimistic. Clearly more work has to be done.

REFERENCES

Bahk, Byong-Hyong, and Michael Gort. "Decomposing Learning by Doing in New Plants," *Journal of Political Economy,* vol. 101 (August 1993), pp. 561–83.

Boskin, Michael J., Ellen R. Dulberger, Robert J. Gordon, Zvi Griliches, and Dale Jorgenson. "Toward a More Accurate Measure of the Cost of Living." Final Report to the Senate Finance Committee from the Advisory Commission to Study the Consumer Price Index, 1996.

David, Paul. "Computer and Dynamo: The Modern Productivity Paradox in a Not-Too-Distant Mirror," in *Technology and Productivity: The Challenge for Economic Policy.* Paris: Organization for Economic Cooperation and Development, 1991, pp. 315–47.

Gordon, Robert J. *The Measurement of Durable Goods Prices.* Chicago: The University of Chicago Press, 1990.

Greenwood, Jeremy, and Boyan Jovanovic. "Accounting for Growth." Manuscript. 1998.

Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell. "Long-Run Implications of Investment-Specific Technological Change," *American Economic Review,* vol. 87 (June 1997), pp. 342–62.

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

Griliches, Zvi. "Productivity, R&D, and the Data Constraint," *American Economic Review,* vol. 84 (March 1994), pp. 1–23.

Hall, G., and S. Howell. "The Experience Curve from the Economist's Perspective," *Strategic Management Journal,* vol. 6 (1985), pp. 197–212.

Hornstein, Andreas. "Reconsidering Investment-Specific Technological Change." Manuscript. 1999.

————, and Per Krusell. "Can Technology Improvements Cause Productivity Slowdowns?" in Ben Bernanke and Julio Rotemberg, eds., *NBER Macroeconomics Annual 1996.* Cambridge, Mass.: MIT Press, 1996, pp. 209–59.

Irwin, Douglas A., and Peter J. Klenow. "Learning-by-Doing Spillovers in the Semiconductor Industry," *Journal of Political Economy,* vol. 102 (December 1994), pp. 1200–27.

Jovanovic, Boyan, and Glenn M. MacDonald. "Competitive Diffusion," *Journal of Political Economy,* vol. 102 (February 1994), pp. 24–52.

Lucas, Robert E., Jr. "Making a Miracle," *Econometrica,* vol. 61 (March 1993), pp. 251–72.

Solow, Robert. "Technical Change and the Aggregate Production Function," *Review of Economics and Statistics,* vol. 70 (1957), pp. 65–94.