Volume Title: NBER Macroeconomics Annual 1996, Volume 11

Volume Author/Editor: Ben S. Bernanke and Julio J. Rotemberg, Editors

Volume Publisher: MIT Press

Volume ISBN: 0-262-02414-4

Volume URL: http://www.nber.org/books/bern96-1

Conference Date: March 8-9, 1996

Publication Date: January 1996

Chapter Title: Can Technology Improvements Cause Productivity Slowdowns?

Chapter Author: Andreas Hornstein, Per Krusell

Chapter URL: http://www.nber.org/chapters/c11030

Chapter pages in book: (p. 209 - 276)

Andreas Hornstein and Per Krusell

UNIVERSITY OF WESTERN ONTARIO, AND FEDERAL RESERVE BANK OF RICHMOND, AND UNIVERSITY OF ROCHESTER, CEPR, AND INSTITUTE FOR INTERNATIONAL ECONOMIC STUDIES

Can Technology Improvements Cause Productivity Slowdowns?

1. *Introduction*

At about the same time as productivity growth slowed down in the United States, some measures of the rate of technological change showed significant strength. In particular, as documented in Gordon (1990), the rate of technology growth embodied in new investment equipment and consumer durables has been remarkable. Gordon's data imply an annual decrease in the price of investment goods in terms of nondurable consumption goods and services of more than 3% on average. A more careful inspection of this relative price series actually also suggests that the rate of technological change was somewhat higher in the late seventies and in the eighties than before. Moreover, using two-digit industry data, McHugh and Lane (1987) study vintage capital effects on productivity and conclude that the rate of capital-embodied technological change went up significantly around the mid-seventies. Although the measurement of technological change is inherently difficult and these findings should be regarded only as suggestive, they do accord with casual observation; for example, the seventies saw the first emergence of robotics and microchip technologies in production processes. The purpose of this paper is to investigate some potentially important implications of rapid technological change for the measurement of the economy's productivity performance.

We thank John Knowles, Max Reid, and Mingwei Yuan for research assistance, and Mark Bils, Robert Gordon, Assar Lindbeck, Torsten Persson, Valerie Ramey, and Julio Rotemberg for comments. We also thank Sam Kortum for providing us with the patent data. Hornstein would like to thank SSHRC for financial support. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Richmond or of the Federal Reserve System.

We focus on two reasons why an increase in the rate of investment-specific technological improvements may lead to a *decrease* in measured productivity performance: learning and quality mismeasurements induced by the increase in the pace of technological progress. First, the adoption of any piece of new capital is associated with learning and initial productivity levels below their full potential. As the rate of technological change increases, relatively more resources are allocated to the new capital, and average knowledge goes down. Temporarily, this produces a slowdown in measured productivity and possibly also a decrease in the output growth rate. Examples of this phenomenon abound; a recent paper by Yorukoglu (1995) emphasizes these effects, and one interesting case study can be found in Gjerding (1991). The kind of technological change which we have witnessed during the last decades can also be argued to represent a change in the way in which capital goods are used in production; there has been a significant move toward labor-saving techniques (robotization), and information technology has inundated the economy. A large management literature argues that as a result of these technological developments, the internal organization of firms has changed substantially; the skill requirements on employees and on management have changed in important ways, especially since many tasks have become computerized. It seems clear that this kind of organizational change itself is an expression of learning; it takes time and resources to reorganize management and the workplace.[1] We formulate a simple model of costly technology adoption which summarizes all the costs of adoption in one variable, and we make a qualitative and preliminary quantitative assessment of the possible magnitude and timing of a productivity slowdown resulting from the increase in the rate of investment-specific technological change that started in the mid-seventies.

Second, there are more general measurement problems associated with the kind of technological change recently witnessed in the United States and other industrialized economies. In particular, there is a widespread perception that the quality component of increases in output (both of intermediate and final goods) is important, and mismeasured quality has been discussed as a potential explanation for the productivity slowdown [see Baily and Gordon (1988) and Griliches (1994) among others]. Similarly, quality improvements have been particularly emphasized in many of the recent contributions to the endogenous-growth literature. In quantitative terms, Robert Gordon's work on the measurement of durable-goods prices (Gordon, 1990) is one of the more striking

1. Lindbeck and Snower (1995) discuss this literature and model the phenomenon of organizational change.

examples of the potential importance of quality improvements. He finds that the adjustment for quality increases the rate of growth of the quantity index for equipment by as much as around 3% annually.

Gordon's focus was on durable goods, which admittedly are prime candidates for goods with important quality improvements, but many other goods have important quality components as well. In fact, large parts of the economy's final output are inherently poorly measured, such as most of the service sector. For example, finance and insurance now produce services well known to incorporate substantial quality components, and these services are often directly tied to the advanced new equipment. Similarly, the retail sector provides hard-to-measure convenience features which to a large extent are made possible by the use of new equipment.

We argue that the higher rate of investment-specific technological change since the early seventies, and structural change toward new kinds of equipment more generally, have also been accompanied by an increase in the quality of output, and hence presumably by an increase in mismeasurement of output. We should therefore expect productivity accounting to result in lower measured performance at around the same time as the productivity slowdown. Moreover, unlike for learning, the slowdown in productivity due to quality mismeasurement will persist as long as the rate of technological progress and/or structural change toward quality does not reverse itself. This indeed is consistent with the data: very few countries have experienced a full recovery from the slowdown, some (including the United States) have had a partial recovery, but a majority have had no recovery. The effects of quality mismeasurements are stronger if the quality component in output has increased in relative terms. Interestingly, we found some indirect evidence from patent data which suggests that quality mismeasurements may have become more important around the mid-seventies.

We develop a simple model where output has both quality and quantity components, and we use this model to show how productivity accounting is affected by different assumptions on structural change, and on what is and is not mismeasured. The exercise serves two purposes. First, our stylized model environment makes conceptually and qualitatively clear what the potential pitfalls of standard productivity accounting procedures are. Second, it attempts a quantitative assessment of how much of the observed productivity slowdown can be attributed to the structural change in investment-specific technology growth.

Our findings are suggestive, but not conclusive. Both learning and quality mismeasurements give rise to slowdowns that are larger in the short run than in the long run. Taken together, they can produce time-series patterns for the slowdown which are not unlike the data. However, the precise patterns and, in particular, the magnitude of the slowdowns depend crucially on parameter values we do not know much about, such as key parameters in the learning technology and the relative importance of quality.

Some additional support for our story can be found in a recent paper by Greenwood and Yorukoglu (1996), who also study the importance of increases in the rate of technological change for productivity slowdowns. Their focus is wholly on learning, and their model of learning is more detailed than ours. In particular, the adoption of new technologies is a choice variable for firms, the learning process is endogenous (it is modeled as requiring skilled labor), and the skill formation is endogenous as well. Moreover, their data analysis includes historical studies of the importance of new equipment during the British industrial revolution in the late eighteenth century and the American industrial revolution in the nineteenth century.[2] The historical data are qualitatively consistent with the story told in their paper as well as with ours: decreases in the price of capital are associated with productivity slowdowns.

Our paper is organized in two parts. In the first part, Section 2, we review the data and the literature on the productivity slowdown. In the second part, Section 3, we conduct our theoretical exercises. First in our data section, we make a review of postwar productivity in the United States and elsewhere, both on an aggregate and on a sectoral level (Section 2.1). In this context, we also discuss the potential importance of mismeasurement by presenting data on the relative importance of the sectors whose output is particularly poorly measured. Next, we go through the implications from adjusting the productivity data using Gordon's price index updates for durable goods (Section 2.2). In Section 2.3, we provide a brief summary of the candidate explanations for the productivity slowdown that have been suggested in the literature. Finally in the data section, we look at some evidence which suggests that the pace of investment-specific technological change has accelerated (Section 2.4). The baseline framework in Section 3 is a simple two-sector model which admits aggregation across sectors. The aggregation allows the learning and the quality mismeasurement hypotheses to be presented in a very simple manner, and we simulate partial models with learning or with quality mismeasurements to illustrate both their qualitative properties and their potential for explaining the magnitude of the observed productivity slowdown. We also simulate a model calibrated to the

2. Their interpretation of the recent slowdown in growth is more specifically tied to information technology than to equipment in general.

United States economy as a quantitative synthesis of the perspective suggested in this paper. Finally, we offer our conclusions in Section 4.

2. *The Productivity Slowdown: Data and Preliminary Assessments*

This section contains a review of productivity statistics in the United States and elsewhere, a very brief summary of the literature on causes of the productivity slowdown, and a section with evidence on a change in the rate of technological change specific to equipment investment. Our data review starts with a description of the postwar development of industry output and productivity in the United States. We then make some international productivity comparisons using aggregate data from a number of developed economies. Since measurement issues are the focus of our theoretical analysis, we also discuss the sectoral productivity measures in the United States and in other countries from the point of view of how well output is measured in the different industries. Explicit and detailed adjustments for quality improvements in durable goods were made in Gordon (1990), and we also discuss the effects of these adjustments on aggregate productivity accounting.

2.1 A REVIEW OF AGGREGATE AND SECTORAL PRODUCTIVITY DATA

Postwar United States data on growth in output and labor productivity are displayed in Table 1 for the years 1954–1993. Industry output is measured in 1987 dollars value added, and labor input is measured by the number of full-time-equivalent employees.[3] Labor productivity is output per unit of labor input. We consider three subperiods: pre-1973, 1973–1979, and post-1979. For the whole period aggregate output growth was 3.1%; it slowed down from 3.7% before 1973 to 2.2% in the mid-seventies and then recovered partially to 2.6%. Similarly, labor productivity growth was 1.3% for the whole period; it was 1.9% before 1973, significantly lower in the mid-seventies (−0.2%), but it had a substantial recovery after 1979 to 1.1%.

Output and productivity growth rates differ substantially across industries during the time period we study. Construction is the industry with the lowest output growth (0.9%) and an average annual decline in labor productivity of 0.5%, whereas the fastest-growing industry is wholesale trade, with an average annual output growth of 4.7% and a labor productivity growth of 2.9%. Although all industries are affected by the 1973–

3. See the Appendix for data sources.

Table 1 OUTPUT AND LABOR PRODUCTIVITY GROWTH, UNITED STATES 1954–1993

	Growth rates (%)							
	Output				*Labor productivity*			
Sector	54–93	54–73	73–79	79–93	54–93	54–73	73–79	79–93
Total private sector	3.1	3.7	2.2	2.6	1.3	1.9	−0.2	1.1
Agric., forestry, fishing	1.3	0.2	−0.0	3.6	1.7	1.9	−1.8	3.0
Mining	1.3	2.5	−3.4	1.8	2.1	3.7	−10.1	5.1
Construction	0.9	1.6	0.3	0.2	−0.5	−0.6	−1.4	0.0
Manufacturing	2.7	3.9	1.4	1.6	2.4	2.8	0.6	2.7
Durables	2.5	3.7	1.0	1.6	2.3	2.4	−0.3	3.2
Nondurables	2.9	4.2	2.0	1.6	2.6	3.4	1.8	1.9
Transport., publ. util.	3.7	4.5	2.8	3.1	2.9	3.9	1.3	2.3
Wholesale trade	4.7	5.2	4.3	4.3	2.9	3.0	1.3	3.3
Retail trade	3.2	3.6	2.1	3.0	0.7	1.1	−1.2	0.9
Finance and insur.	3.7	4.7	4.0	2.3	0.8	1.4	0.3	0.2
Other services	3.9	4.5	3.8	3.1	0.1	0.9	−0.3	−0.8

1979 slowdown in output and labor productivity growth, there are some differences across industries. First, some sectors have dramatic 1973–1979 slowdowns (e.g., labor productivity growth in mining and agriculture fell by 13.8 and 3.7 percentage points, respectively). Second, not all industries recovered from the slowdown after 1979. Agriculture, mining, construction, durable manufacturing, transportation and public utilities, and trade all rebounded partially or fully, whereas nondurable manufacturing, finance and insurance, and other services did not. In particular, while the manufacturing sector as a whole has recovered to its pre-1973 growth rates, the nondurable part of manufacturing has not (and the durable sector has more than recovered).

It is also interesting to note that labor productivity growth rates are not that much lower than output growth rates for many sectors (and higher for some), due to slowly growing or declining employment. However, for the service sector industries, measured labor productivities are substantially lower, reflecting considerable employment growth.

Growth in total-factor productivity (TFP) represents output growth not accounted for by the growth in inputs. Suppose industry i produces output y_i with inputs capital k_i and labor l_i. If production has constant returns to scale and markets are competitive, then the change in TFP between period t and $t + 1$, which is commonly referred to as the "Solow residual," is

$$\Delta \log \mathrm{TFP}_{i,t} = \Delta \log y_{i,t} - a_{i,t} \Delta \log k_{i,t} - (1 - a_{i,t}) \Delta \log l_{i,t},$$

Table 2 TFP GROWTH, UNITED STATES 1954–1993

	Growth rates (%)			
Sector	54–93	54–73	73–79	79–93
Total private sector	0.8	1.3	−0.6	0.7
Agric., forestry, fishing	0.9	−0.1	−2.8	3.8
Mining	−0.1	1.0	−9.1	2.1
Construction	−1.1	−1.7	−2.4	0.2
Manufacturing	1.6	2.1	−0.2	1.6
Durables	1.5	1.7	−0.9	2.2
Nondurables	1.7	2.7	0.7	0.9
Transport., publ. util.	2.4	3.2	0.7	2.0
Wholesale trade	1.6	1.4	0.6	2.3
Retail trade	−0.0	0.2	−1.3	0.2
Finance and insur.	−1.8	−1.6	−1.3	−2.4
Other services	−0.6	−0.6	0.1	−0.8

where $\Delta x_t \equiv x_{t+1} - x_t$ and a_i is capital's share of income in this industry (Solow, 1957). Although this accounting does not tackle the harder question of what determines the growth in inputs and technology, it has proven a very valuable organizing tool for empirical studies of economic growth.[4]

Table 2 shows that the development of TFP in the postwar United States is similar to that of labor productivity.[5] Notable exceptions are finance and insurance and other services: in both these industries output and labor productivity growth slows down in the seventies, but TFP growth improves in that time period. However, both these sectors record significant slowdowns later on. In sum, although capital accumulation does account for a significant fraction of output and labor productivity growth, it does not help us understand the slowdown in the seventies.

Turning now to the international economy, Table 3 shows data from Kendrick (1990). Clearly, the productivity slowdown is worldwide. The slowdown occurs in all of the listed countries, and it is substantial, whether it is measured in labor productivity or TFP growth. The table also shows that the duration of the slowdown differs markedly across countries. Among the 19 countries, 11 have experienced a slowdown which continues unabated throughout the sample period and in some cases becomes significantly worse toward the end of the period (West

4. See for example Jorgenson, Gollop, and Fraumeni (1987) and Denison (1985).
5. For the TFP growth calculations in Table 2 we disaggregate industry capital into equipment and structures. This is useful given our discussion below about the role of equipment-embodied technological change. The Appendix describes our procedure and the sources for the capital stock series.

Table 3 INTERNATIONAL PRODUCTIVITY FACTS, 1960–1988

	Growth rates in percent					
	Labor productivity			TFP		
Country	60–73	73–79	79–88	60–73	73–79	79–88
U.S.	2.8	0.6	1.6	1.8	0.1	0.7
Canada	2.8	1.5	1.5	2.0	0.7	0.3
Japan	9.4	3.2	3.1	6.4	1.8	1.8
Austria	5.8	3.3	1.8	3.4	1.4	0.7
Belgium	5.0	2.8	2.1	3.7	1.5	1.1
Denmark	4.3	2.6	1.5	2.8	1.2	0.8
Finland	5.0	3.4	3.2	3.4	1.7	2.3
France	5.4	3.0	2.4	3.9	1.7	1.5
Germany	4.6	3.4	1.9	2.7	2.0	0.7
Greece	8.8	3.4	0.2	5.8	1.5	−0.7
Italy	6.3	3.0	1.6	4.2	2.2	1.0
Netherlands	4.9	3.3	1.5	3.1	2.0	0.6
Norway	4.1	0.1	2.0	3.6	−0.4	1.4
Spain	6.1	3.8	3.4	4.2	1.7	2.1
Sweden	3.9	1.4	1.6	2.5	0.3	0.9
Switzerland	3.2	0.7	0.9	1.6	−0.9	0.2
U.K.	3.5	1.5	2.6	2.2	0.5	1.9
Australia	3.2	2.0	1.1	2.9	1.2	1.0
New Zealand	1.8	−1.5	1.4	1.0	−2.2	0.6

Germany, Italy, the Netherlands, and Greece). For two countries, the 1979–1988 productivity growth is almost back at its pre-73 level—United Kingdom and New Zealand—but in both cases the pre-73 productivity growth was dismal as well. For the remaining six countries, the United States included, productivity growth has picked up, but it is far from its pre-73 level. It is thus important to note that although the United States productivity slowdown is similar to slowdowns in some other countries, it is among a minority showing partial improvement in the eighties.[6]

2.1.1 Measurement Issues Has the productivity slowdown of the seventies been for real, or does it reflect systematic measurement error? The idea that mismeasurement is potentially important for understanding productivity movements in general and the productivity slowdown in particular is not at all new.[7] Baily and Gordon (1988) discuss some mea-

6. For additional evidence on a recovery of TFP growth in the manufacturing sector in the United States see Gullikson (1995); for more on the absence of a recovery in France, Germany, and the United Kingdom see Lysko (1995).
7. An early reference is Thurow (1981).

surement problems in detail and provide estimates of their importance for the productivity slowdown. More recently, Griliches (1994) has emphasized the importance of poor output measurement for the *unmeasurable sector* of the economy, that is, industries for which output is inherently difficult to measure. Below we document that the increase in relative size of the unmeasurable sector has been significant for a large group of countries in the postwar period. This phenomenon affects our ability to generate reliable productivity measures.

Calculations of real output tend to be less reliable when we do not have a well-defined measure of output. In the absence of such well-defined measures, two methods are usually applied in the United States National Income and Product Accounts (NIPA). The first method extrapolates real output by use of an input series, for example employment.[8] The second method deflates a measure of nominal output by a corresponding consumer price index (CPI). Given the evidence that the CPI overestimates the rate of inflation, as discussed in Shapiro and Wilcox (1996) or in Boskin *et al.* (1995), this means that real output growth is underestimated. These problems are especially relevant when there are substantial changes in the quality of a sector's output. A prime example is the service sector, where there is reason to believe that quality improvements have been substantial.[9]

Griliches (1994) defines the measurable sector to include agriculture, mining, manufacturing, transportation and communications, and public utilities, and the unmeasurable sector to include construction, trade, finance, insurance and real estate, other services, and government. Although there are measurement problems in all industries, we do think this definition offers a reasonable way of illustrating the growing importance of mismeasurement.[10] Table 4 and Figure 1 show the development of TFP in the measurable and unmeasurable sectors. We see that, as of the end of the sample, the unmeasurable sector contributes more than 50% to United States GDP. The table and the figure also seem to say that the distinction between the measurable sector and the unmeasurable sector is a distinction between a "technologically progressive" sector and a "technologically regressive" sector: the unmeasurable sector displays a

8. In the 1991 revision of industry GPO accounts (see De Leeuw, Mohr, and Parker, 1991), this procedure has been replaced when possible by more standard procedures like double deflation. For example, real output in the transportation industry, which used to be based on employment extrapolation, is now calculated using double deflation. There are, however, important exceptions where real output continues to be calculated by extrapolation; one example is the banking sector.

9. For productivity analysis in the service sector see, e.g., Kendrick (1987).

10. For example, Baily and Gordon (1988) point to severe measurement problems in the transportation industry.

Table 4 TFP GROWTH IN THE MEASURABLE AND UNMEASURABLE SECTORS, UNITED STATES 1954–1993

	Growth rate (%)			
Sector	54–93	54–73	73–79	79–93
Total private sector	0.8	1.3	−0.6	0.7
Measurable sector	1.7	2.2	−0.6	2.0
Unmeasurable sector	−0.5	−0.4	−0.7	−0.5

negative time trend for TFP in the postwar period, whereas the measurable sector displays a positive time trend for TFP. For the unmeasurable sector, the slowdown of TFP growth in the seventies meant an even faster decline of TFP. In the period from 1979 on, TFP growth in the measurable sector has essentially recovered, whereas TFP in the unmeasurable sector continues to decline. Of course, it is unlikely that the latter indeed does experience technological regress; the measured decline in TFP may all be due to underestimated output growth.

The unmeasurable sector represents an important component of the economy in most industrialized countries. In Table 5 we report the share

Figure 1 MEASURED TOTAL FACTOR PRODUCTIVITY

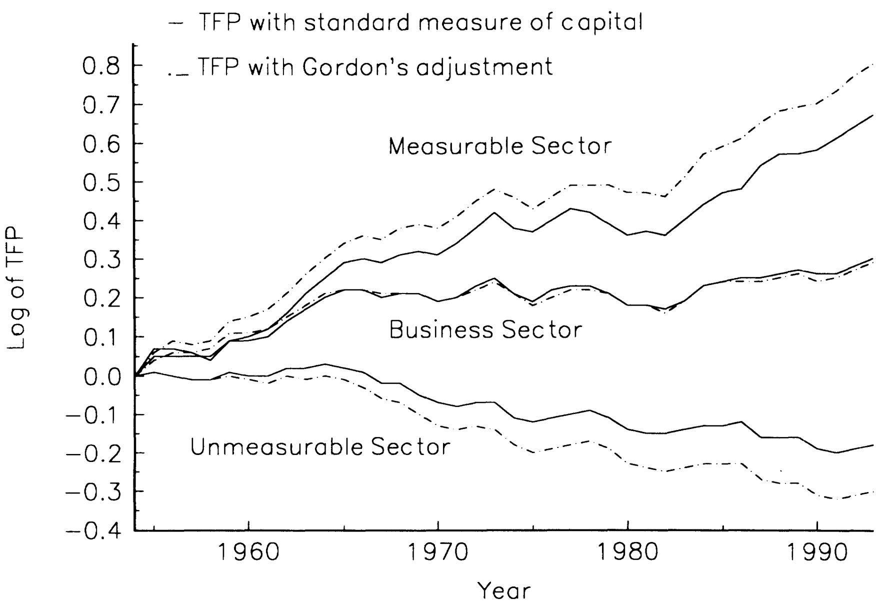

Table 5 THE UNMEASURABLE SECTORS, INTERNATIONAL DATA 1954–1993

	Share of GDP			
Country	1954	1973	1979	1993
U.S.	0.48	0.56	0.58	0.66
Canada	0.43	0.49	0.50	0.61
Japan	0.39	0.49	0.54	0.59
Austria	0.27	0.45	0.48	0.57
Belgium	0.46	0.49	0.54	0.58
Denmark	0.30	0.56	0.56	0.56
Finland	0.29	0.53	0.43	0.48
France	0.36	0.49	0.51	0.61
Germany	0.33	0.43	0.45	0.55
Greece	0.41	0.46	0.49	0.55
Italy	0.37	0.46	0.46	0.59
Luxemburg	0.34	0.44	0.55	0.61
Netherlands	0.39	0.50	0.56	0.62
Norway	0.36	0.43	0.53	0.43
Portugal	0.23	—	0.45	0.48
Spain	0.36	—	0.53	0.61
Sweden	0.46	0.47	0.50	0.58
Turkey	0.31	0.38	0.36	0.42
U.K.	0.41	0.47	0.47	0.58

of the unmeasurable sector's nominal value added in the private business sector's nominal GDP for OECD countries during 1954–1993.[11] Two key facts emerge from Table 5. First, the size of the unmeasurable sector is substantial in all countries; by the end of the sample period, it accounts for more than half of total GDP in most countries, and it accounts for significantly more than half in some. Second, the size of the unmeasurable sector has increased substantially since 1954. This increase occurred in all countries, and it is monotone for the four dates we report in all but a few of the countries.

The facts described above are quite important for the interpretation of aggregate productivity statistics. First, they tell us that measurement problems are important on an aggregate level, since the part of GDP which has fundamental measurement problems is so large. Second, for the postwar period the relative size of the unmeasurable sector, and

11. The data, which are calculated using OECD's National Accounts, Volume II, Detailed Tables, and do not include government in any of the sectors. The beginning year is 1954 for all countries except for Spain, for which it is 1958; the end year is 1991 for Canada, the United States, Luxemburg, Norway, and Spain, 1990 for Belgium, 1989 for Portugal, and 1993 for the rest of the countries.

therefore the measurement problem, have increased substantially. This means that to the extent that the recorded productivity growth rates in the unmeasurable sector are lower than in the rest of the economy, these compositional changes in GDP will themselves lead to a slowdown in aggregate productivity growth.

2.2 QUALITY IMPROVEMENTS IN CAPITAL IN THE POSTWAR PERIOD AND THEIR IMPLICATIONS FOR PRODUCTIVITY ACCOUNTING

One of the few examples of a detailed and systematic assessment of the quantitative importance of quality improvements is the work by Gordon (1990) on durable goods prices. Gordon uses hedonic pricing and other methods aimed at quantifying all the characteristics of new durable goods for evaluating their "true" quantity levels.[12] The results are striking: for durable consumption goods, the rate of price increase is adjusted downward and output growth upward, at an average rate of 1.5 percentage points per year, and for producers' durable equipment (PDE) the corresponding number is 3 percentage points per year. The quality adjustments reflect technological change embodied in new durable goods, so Gordon's work has unveiled important developments in equipment-embodied technological change. In this section we study the implications of Gordon's quality adjustments for the measurement of TFP growth.

2.2.1 Revising the Productivity Accounts Gordon's revision of the durable-goods price series has implications for the measurement of TFP, because it affects our measures of durable-goods industries' outputs and capital inputs used in all industries. First, a higher growth rate of output from durable-goods manufacturing increases labor-productivity and TFP growth in that industry. Second, investment in PDE adds to the stock of capital in all industries, so the upward revision of the growth rate of investment increases the growth rate of the capital stock. This, in turn, leads to a downward adjustment in TFP growth in all industries.[13]

In Table 6 we report the effects on measured TFP growth of adjusting output in the durable-goods manufacturing sector and the equipment stocks in all sectors. For simplicity we have assumed the quality mismea-

12. Except for the case of computers, which are part of the Office Computing and Accounting Machinery (OCAM) category, such methods are not used for the official income statistics.

13. In principle, durable consumption goods represent investment which raises the level of household capital. Although there are good arguments for including household capital and household production into national income accounts, we follow standard procedures and treat durable consumption like any other consumption good.

Table 6 TFP GROWTH, UNITED STATES 1954–1993 BASED ON GORDON'S PRICE ADJUSTMENTS

	Growth rates (%)			
Sector	54–93	54–73	73–79	79–93
Total private sector	0.7	1.3	−0.5	0.6
Agric., forestry, fishing	0.5	−0.5	−3.2	3.4
Mining	−0.7	0.4	−9.7	1.6
Construction	−1.3	−1.8	−2.6	0.1
Manufacturing	2.5	3.2	1.1	2.3
Durables	3.5	4.1	1.7	3.5
Nondurables	1.4	2.4	0.4	0.6
Transport., publ. util.	2.0	2.8	0.3	1.7
Wholesale trade	1.3	1.1	0.3	2.1
Retail trade	−0.3	−0.0	−1.5	−0.1
Finance and insur.	−2.3	−2.1	−1.7	−2.8
Other services	−1.0	−1.1	−0.3	−1.2
Measurable sector	2.1	2.5	0.1	2.3
Unmeasurable sector	−0.8	−0.7	−1.0	−0.8

surement for PDE is the same in all sectors, and we have adjusted sectoral PDE investment and stocks using Gordon's price series. Unfortunately, Gordon's equipment price index is available only until 1983; for the years following 1983 we have used an official price index for the subcategory of equipment where quality is appropriately accounted for and have made an ad hoc, but conservative, adjustment for the other subcategories—a 1.5-percentage-point reduction in the annual growth rate of the official price indexes.[14]

In comparison with the unadjusted numbers in Table 2, we see that the expected changes occur: the TFP growth rates decrease for all sectors except durable-goods manufacturing, by about 0.3–0.4 percentage point on average, and the rate increases for the durable-goods sector by 2 percentage points. The adjustments are larger the larger is the equipment share, so the adjustments in mining, which has the highest capital share, are the largest outside durable manufacturing. However, the adjustments here have no new implications for the period of the slowdown: the adjustments are quite uniform over the whole period. Also note that we have plotted the adjusted series for the measurable and

14. This procedure was suggested to us by Robert Gordon. More elaborate procedures based on estimating the patterns of adjustments made by Gordon and forecasting the quality adjustments after 1983 have been tried elsewhere (see Krusell, et al., 1995) and result in more drastic adjustments.

unmeasurable aggregates in Figure 1. The adjustments increase overall TFP growth in the measurable sector and decrease it in the unmeasurable sector (durable manufacturing is regarded as measurable).

2.2.2 A Consumption–Investment Breakdown of Productivity Improvements
Investment-specific technological change, that is, technological improvements in the sectors producing investment goods, is different in nature from that in other sectors. Unlike technological improvements in other sectors, investment-specific technological change will increase labor productivity and output growth in all sectors, since capital is used as an input in all sectors. For this reason, it is instructive to take a slightly different perspective on productivity accounting than the one adopted above.

Suppose we think of the economy as consisting of two sectors, one producing consumption goods c and one producing investment goods i. Each good is produced using capital and labor as inputs:

$$c_t = \gamma_{c,t} f_c(k_{c,t}, l_{c,t}),$$
$$i_t = \gamma_{i,t} f_i(k_{i,t}, l_{i,t}).$$

Let k_t and l_t be the economy wide endowments of capital and labor at time t, and suppose that inputs are freely mobile between sectors. Furthermore assume that production has constant returns to scale and that isoquants have the same shape in the two sectors, i.e., that $f \equiv f_c = f_i$. Then perfect competition in all markets implies that total output in terms of consumption goods can be written as follows:

$$c_t + \frac{\gamma_{c,t}}{\gamma_{i,t}} i_t = \gamma_{c,t} f(k_t, l_t).$$

The reciprocal of the ratio $q \equiv \gamma_i / \gamma_c$ represents the relative price of investment goods. We identify the investment-specific component of productivity with the ratio q. Alternatively, this ratio could be called capital-embodied productivity, since, relative to consumption, the growth rate of q will equal the rate at which capital goods production becomes more efficient over time.[15] Similarly, we may denote by $z \equiv \gamma_c$ the sector-neutral productivity.

Measuring total output in units of consumption goods and identifying output with an aggregate production function means that aggregate

15. To see this, define investment in consumption units $\hat{\imath}$ and note that the amount of new capital produced by this investment equals $i = \hat{\imath}q$.

productivity should be interpreted as the economy's ability to produce consumption goods, given values for total inputs. Similarly, we can define output in terms of investment goods, and measure the economy's ability to produce investment goods for given total inputs. This amounts to multiplying the above constraint with γ_i/γ_c, which yields an output definition which is associated with the productivity level γ_i. In other words, "aggregate productivity" defined this way really is sectoral productivity, and the aggregate productive ability of the economy is best described by a vector (γ_i, γ_c). Thus, the additional theoretical assumptions allow us to recover the sector-specific productivity parameters without the knowledge of sector-specific input levels.

Equipped with a standard capital accumulation equation and some specification of savings behavior, we now have a version of the standard neoclassical growth model which allows for investment-specific technological change. As such, this framework is well suited for analyzing the relative importance of investment-specific technological change for long-run output growth. Such an analysis is performed in Greenwood, Hercowitz, and Krusell (1995), and it proceeds in two steps.

First, the growth rate of q is identified with the rate of decline of the relative price of capital goods, as given by Gordon's price deflator for PDE divided by the price deflator for nondurable consumption goods and services.[16] The updated version of this series is plotted in Figure 2. According to Figure 2, the relative price of PDE declines (that is, investment-specific technological knowledge increases) at an annual rate of about 3%. Also, it appears that the rate of price decline is larger towards the end of the sample, and that the change in trend occurs in the mid-seventies. We will get back to this issue shortly.

Second, the growth rate of this ratio over the sample is used to calculate how the long-run growth rate which follows from the balanced-growth path of the model depends on the two kinds of investment-specific and neutral technological change, respectively. Greenwood, Hercowitz, and Krusell (1995) find that the former accounts for around 60% of total consumption growth, which is what is of importance to consumers living in this kind of economy. As part of the procedure, a series for the neutral technological change z is obtained. The series is displayed in Figure 3.

This graph shows a drastic version of the productivity slowdown: neutral productivity increases until the mid-seventies, after which it falls uninterruptedly until the very end of the sample, when it increases again. There are two reasons for the drastic productivity slowdown/fall. The first reason is the increase in the rate of growth of the capital stock

16. That analysis also distinguishes between investment in PDE and in structures.

Figure 2 PRICE OF PDE RELATIVE TO NONDURABLE CONSUMPTION AND SERVICES

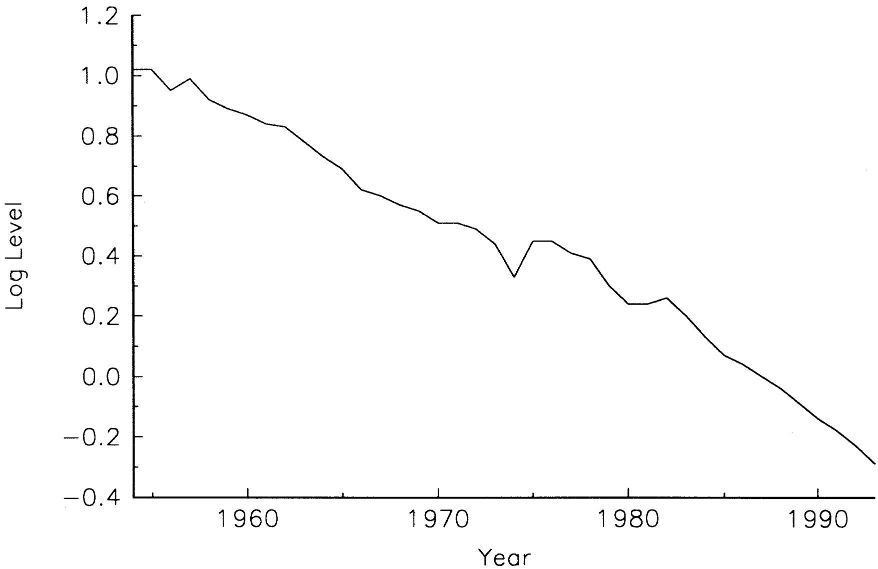

Figure 3 NEUTRAL TECHNOLOGICAL CHANGE

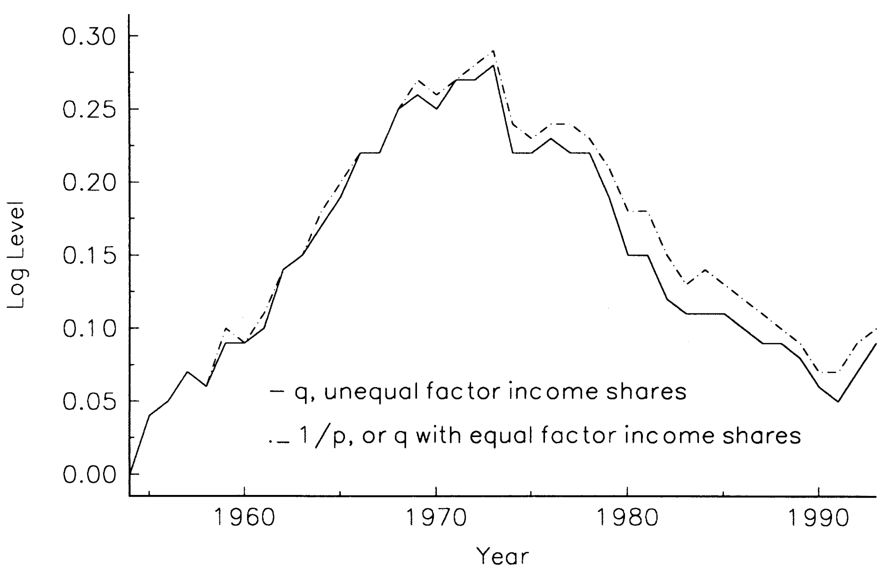

implied by Gordon's revisions of the relative efficiency of new investment. Second, we measure the economy's current ability to produce nondurable consumption goods and services, but these goods are to a large extent supplied by the unmeasurable sector. We have seen above that TFP in that sector actually declined in the postwar period, and thus our result is perfectly consistent with the sectoral revisions above. To evaluate the economy's ability to transform current consumption into future consumption, the growth of q becomes relevant.[17]

This paper examines possible implications of investment-specific technological change for productivity accounting. In particular, we use theoretical model economies to examine whether this phenomenon, together with learning about new technologies and mismeasurement of the quality improvements in new goods, can help us understand the recorded productivity slowdown. Before we proceed to that analysis, however, we have two more topics to cover. As a background, we will first provide a brief review of explanations for the productivity slowdown. Second, we take another look at technological change occurring in the investment-goods sector and provide some additional evidence that the rate of investment-specific technological change has increased sometime during the mid-seventies.

2.3 EXPLANATIONS FOR THE SLOWDOWN: A BRIEF LITERATURE REVIEW

A variety of explanations have been proposed to account for the productivity slowdown. At this point we provide a brief review of the more popular ones. We would like to suggest that while these explanations account for some of the observed slowdown in TFP growth, the larger part of the slowdown remains unexplained. The review will be quite brief and sketchy. For a more comprehensive treatment the interested reader may consult any one of the many excellent survey papers, summaries, and conferences volumes on the subject.[18]

A reasonably comprehensive list of potential explanatory factors is as follows. We exclude the "measurement explanation" here, since it has already been discussed and will be elaborated on in later sections.

Decreases in labor quality. In the productivity calculations above, no adjustments are made for changes in labor quality. Many authors have

17. Note also that the ability to produce durable consumption goods is not directly given by the graph, but needs to be adjusted by the relative price increase for nondurables in terms of durables.

18. See, for example, Cullison (1989), the 1988:2(4) issue of the *Journal of Economic Perspectives*, and Baily and Gordon (1988).

used educational attainments, possibly corrected for the quality of the educational achievements (e.g., as measured by SAT test scores), and other aspects of changes in the heterogeneity of the labor force (the age–race–sex distribution) to revise the labor input. The difficulty with this explanation is that no substantial changes in the labor input can be detected at the time the slowdown began. In addition, many of the changes in the quality of the labor force are specific to the United States and therefore do not explain the international slowdown. Views on the importance of this explanation, however, differ.[19]

The oil shock(s). To many, the most attractive explanation for the slowdown is the oil shock, because the timing is right and because it is common to all countries.[20] There are problems with this explanation. Taking into account the cost share of energy, and the modest slowdown that occurred in actual energy input, only a very small fraction of the slowdown can be accounted for. Indirect effects through capital obsolescence are also problematic: given the small energy cost share, the massive move to alternative kinds of capital which would be needed to motivate the large slowdown seems to contradict profit maximization. In addition, we have not observed the changes in used-capital prices that would follow from the obsolescence explanation. Furthermore, following the reduction in real oil prices in the eighties we have not observed an increase in TFP growth rates which is comparable in magnitude to the slowdown in the seventies. All the same, the timing of this explanation is "too good to be a coincidence" in the views of many, and other indirect, but not spelled-out effects, such as macroeconomic disruptions, have been mentioned.[21]

A slowdown in R&D and the number of technological innovations. In the United States, the R&D share of total expenditures declined in the mid-sixties, and the number of patents per R&D dollar has also declined. However, the decline in R&D expenditure is specific to the United States, and R&D expenditures have increased again without a concomitant increase in productivity. Moreover, the R&D explanation is tailored to the manufacturing industry, and does little to address the decline in service-sector productivity. Did the number of new innovations go down, and are we experiencing diminishing returns to R&D and "technological exhaustion"? First, in the United States the decline

19. For references, see Denison (1985), Darby (1984), Baily (1981a), Baily and Gordon (1988), Jorgenson *et al.* (1987), and Dean, Kunze, and Rosenblum (1988).

20. Baily and Gordon (1988) do provide an argument why the timing is not perfect: a slowdown had already begun in several sectors before the oil shock.

21. For references, see Nordhaus (1982), Summers (1982), Baily (1981b), Jorgenson *et al.* (1987), Bruno (1981), Bruno and Sachs (1982), Hulten, Robertson, and Wykoff (1989), and Olson (1988).

in the number of granted patents can be explained by budget cuts and a decline in resources allocated to patent granting. Second, in our view, it seems difficult to argue that the last two decades have been characterized by particularly slow technological change, considering the rapid expansion in information technology and other high-technology areas.[22]

Regulations, cultural change, labor disputes, management failures. It has been argued that the increase in the number and strictness of regulations in the United States during the second half of the postwar period may have played an important role in lowering productivity. Similarly, there have been increases in crime rates and labor-market disruptions which have the potential to lower productivity. Management failures which could also be reflected in a decrease of measured productivity have also been stressed. Although there are merits to all these explanations, they have the usual problems: there is no perceived sharp increase in any one of these factors in the mid-seventies, and although some of these problems did occur in some other countries, none of them is worldwide.[23]

The coincidence of a number of sector-specific problems. One approach is to analyze what might have caused the slowdown sector by sector. For example, it has been argued that the slowdown in construction is due to unionization, as well as to specific problems with output deflation. Mining has had problems because marginal costs of extraction have risen rapidly, especially during the time the oil price and production increased. The electric utility industry is characterized by large fixed costs and very low marginal costs, so when demand decreases, as it did after the oil shocks, measured productivity falls substantially. Although some of the sector-specific explanations have common causes, such as the oil shocks, a more complete analysis of all sectors is unlikely to satisfy the timing requirement and to be valid for other countries.[24]

In conclusion, views differ widely on the quantitative importance of the different explanations for the slowdown. Denison (1985) can account

22. For surveys and case studies on R&D, see Griliches (1988, 1994) and Baily and Chakrabarti (1988). For the technology exhaustion hypothesis, see Baumol and Wolff (1979), and Nordhaus (1982).
23. For references on the effects of regulations, see Denison (1985), Norsworthy, Harper, and Kunze (1979), and Christiansen and Haveman (1981). For cultural aspects, see Denison (1985), and Naples (1988). For labor-market disruptions, see Denison (1985), Gordon (1981), and Naples (1981), and for management failures see Hayes and Abernathy (1980), Dertouzos, Lester, and Solow (1990), and Summers (1982).
24. For references, see Baily and Gordon (1988), Allen (1985), and Thurow (1987).

for about one-third of the slowdown in the seventies with a subset of the explanations listed above. Others claim greater success, but it seems fair to say that not more than half of the slowdown has been accounted for.

2.4 EVIDENCE ON STRUCTURAL CHANGE

We propose an alternative explanation of the productivity slowdown: we suggest that the measured decline in productivity growth can in part be attributed to an increase in the rate of investment-specific technological change. In this section we present evidence for an increase in the rate of investment-specific technological change during the seventies. In the next sections, we will then discuss why an increase in the rate of technological change may lead to a decrease in measured productivity growth. Because technologies are available worldwide, our explanation can account for the simultaneous decline in productivity growth among industrialized countries.

The task of quantifying the rate of technological change, not to mention detecting a long-term change in this rate, is difficult. There is ample anecdotal evidence on important technological improvements, most of them capital-embodied, which have occurred during the last decades. Many of these improvements have been associated with the introduction of microprocessors and computers into the production process. Computers have made possible new organizational structures, and they have been incorporated into other capital goods. In manufacturing, numerically controlled machine tools, robotization, and automatic assembly have been introduced in many production processes [see Edquist and Jacobsson (1988) for a discussion]. Faster and more efficient means of telecommunication and transportation have also been developed. It is of course difficult to date any of these developments precisely, but many of them did appear in the seventies. Of course, the critical reader should then note that the fifties and sixties also saw many advances in the production of consumer electronics, cars, and so on. Although most of the anecdotal evidence which we have encountered for the earlier period is less equipment-related than for the period of the slowdown and thus not really contradictory with our thesis, it is clear that we need to go beyond speculation about an increase in the rate of capital-embodied technological change based purely on anecdotes. Therefore, we investigate two measures of the aggregate rate of investment-specific technological change. These measures are based on different kinds of data, and thus complementary. They do speak in favor of a structural break in the growth rate of capital-embodied technology.

2.4.1 Use of the Relative-Price Data Hulten (1992) and Greenwood, Hercowitz, and Krusell (1995) identify the growth rate of capital-embodied technological change with the rate of decline in the relative price of investment goods. Their procedure, as outlined in Section 2.2.2, relies on assumptions similar to those in Solow (1957). The relative price of PDE based on Gordon's price series is displayed in Figure 2. Inspection of this figure suggests that starting in the mid-seventies the relative price of new capital declined at a higher rate, about one percentage point more on an annual basis. This would indicate an accelerated rate of capital-embodied technological change. Before we test for a structural break in the relative-price series, we want to discuss in more detail the identification of the relative price of capital with capital-embodied technological change.

A change in a relative price can reflect a change in relative productivity, or it may simply reflect substitution in production. The analysis in Section 2.2.2 shows that the relative price of investment goods reflects the relative productivity of the investment sector only if production is competitive, inputs are mobile across sectors, and the production isoquants are the same in the investment- and the consumption-goods sector. In terms of a Cobb–Douglas production function, the last condition means that the capital and labor income-share parameters have to be the same. As we have pointed out earlier, income shares differ across sectors; in particular, the sector producing durable goods has one of the lowest capital income shares. Greenwood, Hercowitz, and Krusell (1995) show that on the balanced-growth path the behavior of the relative price of investment goods depends crucially on the relative income shares in the consumption- and investment-goods sectors. In particular, if the capital income share is lower in the investment sector and there is no investment-specific technological change, we should observe an increase in the relative price of investment goods. Since we observe the opposite, the decline in the relative price of investment goods must *under*estimate the growth rate of investment-specific technological change. More to the point, to the extent that income shares in the two sectors are different and trend differently over time, any inference about the rate of technological change which assumes constant and equal shares will be misleading.

To be more concrete, assume that production in the investment- and consumption-goods sector is Cobb–Douglas,

$$c_t = z_t k_{c,t}^{\alpha_{c,t}} l_{c,t}^{1-\alpha_{c,t}},$$
$$i_t = q_t z_t k_{i,t}^{\alpha_{i,t}} l_{i,t}^{1-\alpha_{i,t}},$$

and that income shares are not constant. Assume perfect competition in production, and let p_t denote the relative price of capital. It is easy to show that the relative productivity of the investment-goods sector q_t satisfies the following relationship when inputs are freely mobile across sectors:

$$\frac{1}{q_t} = \frac{c_t}{i_t} \frac{k_{i,t}^{\alpha_{i,t}} l_{i,t}^{1-\alpha_{i,t}}}{k_{c,t}^{\alpha_{c,t}} l_{c,t}^{1-\alpha_{c,t}}},$$

where $k_t = k_{c,t} + k_{i,t}$ and $l_t = l_{c,t} + l_{i,t}$ and

$$l_{ct} = \frac{l_t}{1 + \frac{p_t i_t}{c_t} \frac{1 - \alpha_{i,t}}{1 - \alpha_{c,t}}}$$

and

$$k_{ct} = \frac{k_t}{1 + \frac{1 - \alpha_{c,t}}{1 - \alpha_{i,t}} \frac{\alpha_{i,t}}{\alpha_{c,t}} \frac{l_{i,t}}{l_{c,t}}}.$$

Using data on aggregate quantities $\{k_t, l_t, c_t, i_t\}$, prices $\{p_t\}$, and income shares $\{\alpha_{c,t}, \alpha_{i,t}\}$, we can use these equations to construct a series for the relative productivity of the investment sector, q_t. In particular, we use time-series data on the capital income share to isolate the effect of any trend change in this variable on the relative price.[25] In the postwar United States the capital income share in the durable-goods manufacturing sector has declined relative to the share in other sectors. Following our argument above, everything else equal this should lead to an *increase* in the relative price of equipment. Hence, our adjustment procedure will imply an increasing trend in the rate of investment-specific technological change. The reciprocal of the relative price p_t and our measure of investment-specific technological change q_t are graphed in Figure 4.

The price of PDE relative to the price of nondurable consumption and services declined over the postwar period, and the rate of this decline increased in the mid-seventies. A simple regression of the relative price of PDE on a time trend and a change in trend in 1973 shows that there was a

25. In the actual implementation, we have used a slightly more elaborate setup which uses both equipment and structures, and the investment sector thus represents the sector producing equipment. For a more detailed description of the data and the procedure used, see the Appendix.

Figure 4 INVESTMENT-SPECIFIC TECHNOLOGICAL CHANGE

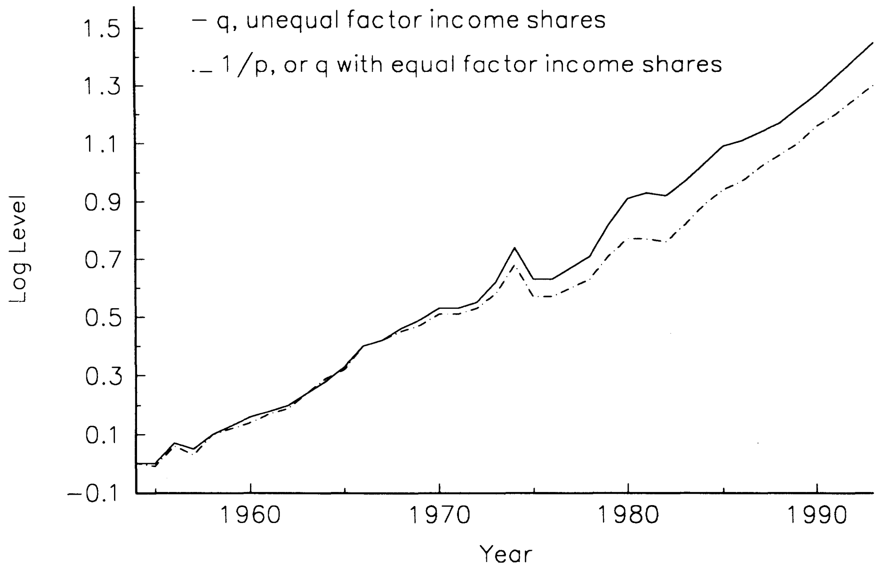

statistically significant increase in the rate of decline from 2.9% before 1973 to 3.6% after 1973; see column 1 of Table 7. In the second column of Table 7 we present the results from a regression of the derived inverse relative productivity series on a linear time trend with a change in trend in 1973. The derived inverse relative productivity series exhibits an even faster rate of decline, and the rate of decline increases from 3.2% before 1973 to 4% after 1973. The change in the rate of decline is highly significant.

Finally, we should comment on the assumption that prices are given by marginal productivities. If markups are variable, changes in relative prices need not reflect relative productivity change. It is thus possible that the decrease in the rate of decline of the relative price of equipment reflects a decline in markups. This is consistent with the notion that we have seen an increase in international competition for at least a subset of the economy's products. However, note that for this alternative explanation of the structural change in the relative-price time series, two elements are necessary. First, it is necessary that the decline in markups have the right time-series pattern. In particular, to explain the change in the trend of the relative price around the mid-seventies, the markup would need to have started to fall around the same time, and it would need to have *continued to fall* throughout the rest of the period (actually, the fall in the relative price even seems to accelerate toward the end of

Table 7 STRUCTURAL CHANGE

Dependent variable	$\log p_t$	$\log (1/q_t)$
Constant	0.0415	0.0569
	(0.0204)[a]	(0.0169)[a]
Trend	−0.0289	−0.0322
	(0.0014)[b]	(0.0012)[b]
Trend change for $t \geq 1973$	−0.0058	−0.0085
	(0.0026)[a]	(0.0021)[b]

[a]significant at the 5% level.
[b]significant at the 1% level.

the period). Second, it is necessary for the decline in markups to have been larger for equipment goods than for consumption goods, since otherwise there would be no change in the *relative* price which we focus on. Although we do not suggest to rule out the declining-markup explanation, we do not know of any evidence supporting the necessary elements for this explanation.

2.4.2 Other Evidence on Capital-Embodied Technological Change The literature on vintage-specific productivity effects can also be used to shed light on the rate of capital-embodied technological progress. In fact, McHugh and Lane (1987) looked at precisely the issue we are interested in. They study the effect of the age of capital on labor productivity using data from two-digit manufacturing industries in the United States. Using a framework which builds on Solow (1959), they conclude that (1) labor productivity declines significantly with increasing age of the capital stock, and (2) the negative effect of the age of capital is significantly stronger for capital installed after 1974. That is, a one-year difference in capital vintage corresponds to a larger productivity difference if the capital was installed after 1974. McHugh and Lane hence conclude that there was an increase in the technological advancement embodied in capital.

The results of McHugh and Lane are derived from a model structure and a concept of technological change quite like ours, and they are consistent with our findings. Their results provide additional, independent evidence for our hypothesis of an increase in the rate of investment-specific technological change sometime in the mid-seventies.

3. *Theoretical Framework and Analysis*

The increase in the rate of investment-specific technological change can reduce measured productivity growth for several reasons. A higher rate of technological change means that new technologies with which producers have less experience are introduced at a faster rate. This can lead to temporarily lower output growth. Moreover, it causes problems for productivity measurement, since experience is unobserved. Faster technological change can also mean that new kinds of goods which differ substantially in their characteristics from existing goods are introduced at a faster rate. This makes the measurement of output more difficult.

There is a substantial body of research which shows that learning, in particular learning by doing, has important effects on productivity. Over time individuals learn how to perform certain tasks, production sites become more efficient, and productivity increases. Learning curves which relate productivity to some measure of accumulated experience have been estimated for a large number of applications: well-known examples in economics include Rapping (1965) on shipyards and Alchian (1963) on airframe production.[26] While there is agreement on the fact that there is learning, there are few explicit models of the learning technology itself.[27] In our work we will simply assume that learning about new technologies is necessary and that it is exogenous. We will study the quantitative implications of learning for the measurement of productivity growth in a simple vintage model of growth.

In previous sections we have argued that for a large part of the economy measures of output are not reliable. We will develop a simple model in which we can discuss the problem of mismeasured output and how it relates to an accelerated pace of technological change. Again, the purpose is to quantify the implications for the slowdown of productivity growth in a simple vintage model. In a final section, we bring the two explanations of the productivity slowdown together, and we use actual United States input and relative price series for a quantitative evaluation.

3.1 SLOWDOWNS DUE TO LEARNING

We now turn to productivity measurement in a growth model when there is learning about new plants or new capital goods. Suppose that any investment in period t is incorporated into a "vintage t" plant. In other words, we do not consider "retooling." In each existing plant learning proceeds at an exogenous rate. The production function with learning thus reads

$$y = \gamma T k^{\alpha} l^{1-\alpha},$$

26. For a survey on learning by doing in the management literature see Yelle (1979).
27. Recently Jovanovic and Nyarko (1995) have interpreted estimated learning curves within a Bayesian learning model.

where T represents what has been learned about the plant, and T is increasing with time. It is possible to think of the structure in this section as the consumption–investment two-sector model of Section 2.2.2 with γ here representing sector-neutral technology (γ_c), and so on.

3.1.1 A Vintage Model with Aggregation At time t, total investment is i_t, and it is all put into a new plant (or many smaller plants). With the assumption that each unit of capital depreciates at rate δ per period, the output at time t of a plant which was set up with investment in $t - \tau$, which we label $y_{t,\tau}$, is

$$y_{t,\tau} = \gamma_t T_{t,\tau} \left[i_{t-\tau}(1 - \delta)^{\tau-1} \right]^{\alpha} l_{t,\tau}^{1-\alpha}$$

for all t and all $\tau \geq 1$, where $T_{t,\tau}$ is the experience level and $l_{t,\tau}$ the amount of labor used at this plant. Total production of goods is

$$y_t = \sum_{\tau=1}^{\infty} y_{t,\tau}.$$

When labor can be allocated freely across vintages at any moment in time, optimal allocations satisfy

$$l_{t,\tau} = l_{t,1} r_{t,\tau},$$

where

$$r_{t,\tau} = \frac{i_{t-\tau}}{i_{t-\tau}} (1 - \delta)^{\tau-1} \left(\frac{T_{t,\tau}}{T_{t,1}} \right)^{1/\alpha}.$$

This condition follows from equalizing the marginal product of labor across plants. The vector $(1, r_{t,2}, r_{t,3}, \ldots)$ thus describes the relative allocation of labor across vintages and, in particular, across learning levels. Using the resource constraint for labor, it is straightforward to show that

$$l_{t,1} = \frac{l_t}{\sum_{\tau=1}^{\infty} r_{\tau,\tau}},$$

where l_t is the total amount of labor input at time t. Similarly, total output at time t satisfies

$$y_t = \gamma_t \tilde{k}_t^{\alpha} l_t^{1-\alpha},$$

where

$$\tilde{k}_t = \sum_{\tau=1}^{\infty} i_{t-\tau}(1-\delta)^{\tau-1}T_{t,\tau}^{1/\alpha}.$$

Hence, this economy admits sectoral aggregation: there exists an aggregate capital measure $\tilde{k}_t$, defined as above, such that output at any point in time is given by a function of this capital measure, total labor input, and the productivity parameter alone.

It is clear from the above that if there were accurate measures of the learning parameters $T_{t,\tau}$, aggregate growth accounting would proceed as in the previous section, replacing the total capital stock k_t with $\tilde{k}_t$, and the TFP growth rate would be

$$\Delta \log z_t = \Delta \log y_t - \alpha \Delta \log \tilde{k}_t - (1 - \alpha) \Delta \log l_t.$$

This procedure would indeed lead to accurate measurement of technology change: $\Delta \log z_t = \Delta \log \gamma_t$. Instead, however, we assume here that capital is measured as in the national accounts, which do not adjust for learning levels. Thus, the measured capital, $\hat{k}_t$, is calculated using $\hat{k}_{t+1} = (1 - \delta)\hat{k}_t + i_t$, so that

$$\hat{k}_t = \sum_{\tau=1}^{\infty} i_{t-\tau}(1 - \delta)^{\tau-1},$$

and growth accounting results in measured TFP growth

$$\Delta \log \hat{z}_t = \Delta \log y_t - \alpha\Delta \log \hat{k}_t - (1 - \alpha) \Delta \log l_t.$$

For this analysis, we assume that output and labor input are well measured. Given that learning levels are not well measured, however, there is no reason to expect that $\Delta \log \hat{z}_t = \Delta \log z_t$ in general. On a balanced growth path, however, the growth rate of productivity is accurately measured. Specifically, suppose that $\gamma_t = \gamma^t$ and that investment-specific technological change is given by $\gamma_{qt} = \gamma_q^t$, so that output in consumption units grows at $\gamma^{1/(1-\alpha)}\gamma_q^{\alpha/(1-\alpha)}$ and investment and capital at $(\gamma\gamma_q)^{1/(1-\alpha)}$. It is easy to see that in this case the vector $r_{t,\tau}$ has to be constant over time, independently of the assumed learning process. Hence, the discrepancy between $\hat{k}_t$ and $\tilde{k}_t$ is constant over time, and it follows that $\Delta \log \hat{k}_t = \Delta \log \tilde{k}_t$.

In the following experiment, we represent structural change by a one-time permanent change in the growth rate of γ_{qt}. More specifically, we assume that the economy is on a balanced growth path with both sources of productivity growing at constant rates, but that the rate of investment-specific technological change increases once and for all, whereas sector-neutral technological change stays the same. For this experiment, we use the actual pre- and post-1973 growth rates calculated in Section 2.4. This experiment will lead to changes in the measured growth rate of TFP, $\Delta \log \hat{z}$, even though the true growth rate is unchanged. The mismeasurement is temporary: as the economy converges to the new balanced-growth path, the error converges to zero.

The parameters are calibrated as follows. We choose $\alpha = 0.3$, $\delta = 0.1$, an investment rate of 0.12, and $\gamma = 1.01$. The initial value for γ_q is taken to be 1.032, and its new value 1.041. The learning technology is specified as

$$T_{t,\tau} = 1 - \lambda^{\tau-1}(1 - T_{t,1}),$$

that is, learning is geometric with convergence rate λ from a starting value of $T_{t,1}$. We consider two ways of selecting the starting values.

First, we look at constant starting values. Following the study by Bahk and Gort (1993) of the long-term experience accumulation in new production plants, we chose $\lambda = 0.7$ and $T_{t,1} = 0.8$. Compared with most of the empirical literature on learning, these values imply a relatively slow learning rate and a small scope of learning. This literature, however, has typically been focused on well-defined learning tasks and not dealt with the kind of complex learning that is a likely result of the technological change we consider here. We actually consider our calibration conservative, since for example the organizational changes in the workplace implied by the availability of information technology (IT) seem more complex and costly than the learning processes analyzed in Bahk and Gort (1993). Moreover, it is arguable that the new information and telecommunication technologies introduced since the seventies have incorporated a new kind of learning or adjustment element because of network externalities: firms do not benefit from and cannot fully learn about their new investment goods until other firms invest as well.

Second, we consider the starting value to be a function of the rate of investment-specific technological progress. In the context of IT investment, Yorukoglu (1995) argues that there are important compatibility problems across different types of capital. In particular he suggests that the more advanced the new equipment is relative to existing equipment, the lower is the initial experience with the new equipment. We consider it reasonable to adopt this approach also when capital is defined more broadly. To simplify things we assume that at the time the rate of investment-specific technological change increases, the starting value for experience declines to $T_{t,1} = 0.6$.

Figure 5 LEARNING

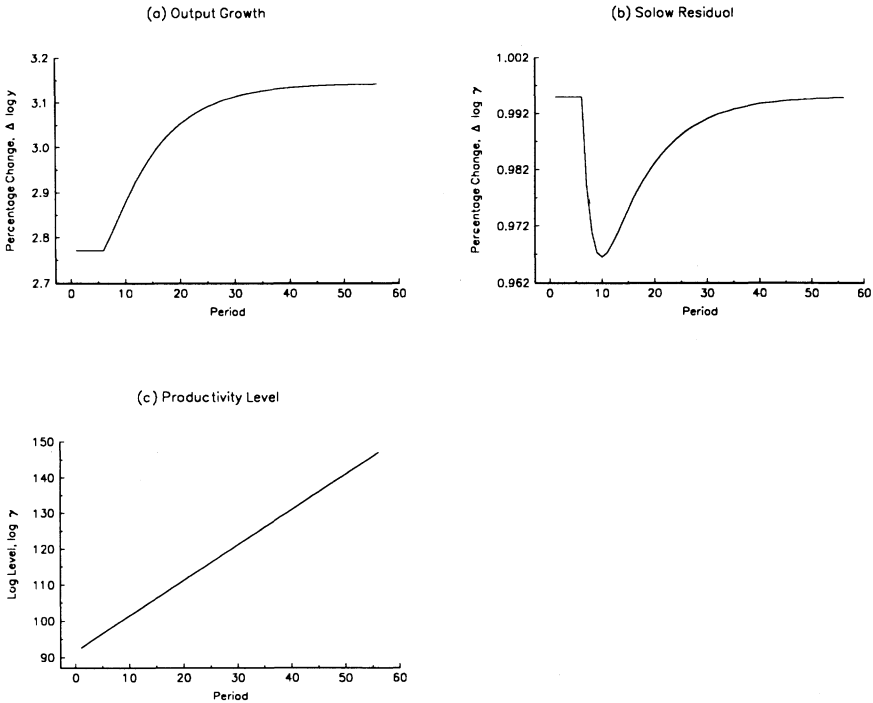

3.1.2 Results Our results for the two cases are displayed in Figures 5 and 6. Figure 5 assumes no compatibility problem, whereas Figure 6 does. Common to these figures is an initial slowdown and a subsequent recovery of measured TFP growth, $\Delta\hat{z}$. Note in Figure 5, however, that for the learning parameters we selected, there is no slowdown in output growth. Since total employment is fixed, this also means that there is no slowdown in labor productivity growth. Only with a compatibility problem do we observe a slowdown in output growth. The slowdown in productivity growth reflects reallocation of labor toward more recent vintages due to the higher rate of technological progress, and with more labor concentrated in recent vintages, average learning factors necessarily drop. Notice also that with a compatibility problem the decrease in the average learning level among plants causes a permanent level drop in measured TFP, even though the TFP growth rate comes back to its true value. The model with a compatibility problem does produce a noticeable slowdown: measured TFP growth declines by $1\frac{1}{2}$ percentage points. However, the slowdown is short-lived: it lasts no longer than 5 years.

Figure 6 LEARNING WITH A COMPATIBILITY PROBLEM

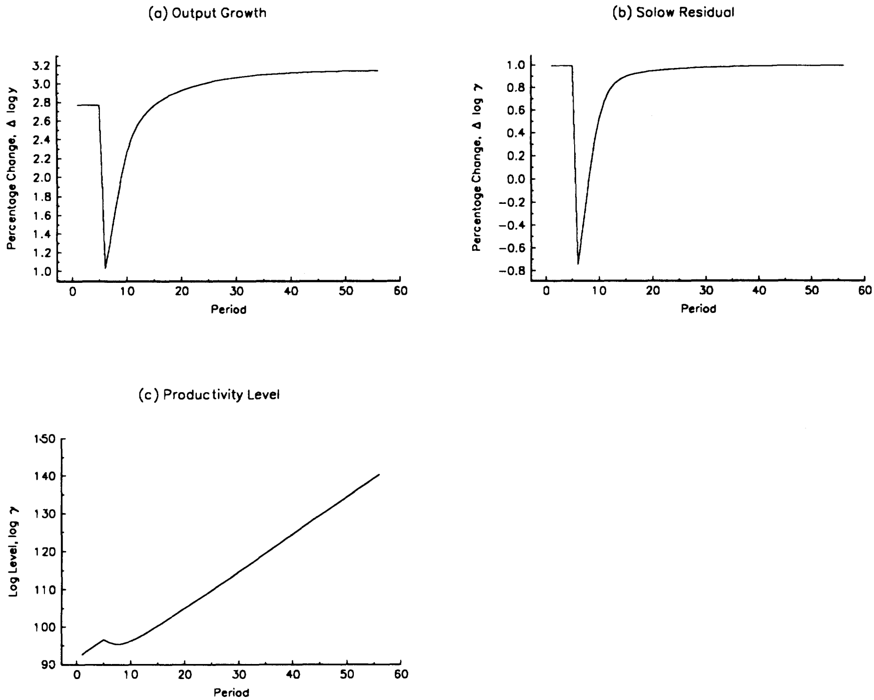

We have also used our learning formulation to adjust actual United States capital stock data by sector, which allows a revision of the TFP figures. The revisions, which are based on the same parameters as Figure 6 and which also use capital stocks based on Gordon's price-index data, are displayed in Table 8. The adjustment of the capital stock for learning increases measured TFP growth. The effect is not particularly large overall, on average about 0.2 percentage points for the 1954–1993 period. For the productivity slowdown of the mid-seventies, the effect is more substantial. Overall measured TFP growth for the 1973–1979 period is about 0.7 percentage points higher than without an adjustment for learning, with the most dramatic effect on finance and insurance, where TFP growth is now about 1.2 percentage points higher. We also include numbers for the measurable and unmeasurable sectors as defined in Section 2.1. The adjustment of the capital stock for learning affects measured TFP about the same way in the measurable and unmeasurable sector. Learning alone has a small effect, increasing measured TFP growth by about 0.1

Table 8 TFP GROWTH, 1954–1993 WITH ADJUSTMENTS FOR LEARNING

	Growth rate (%)			
Sector	54–93	54–73	73–79	79–93
Total private sector	0.9	1.4	0.2	0.5
Agric., forestry, fishing	0.6	−0.4	−2.4	3.3
Mining	−0.7	0.5	−8.4	1.1
Construction	−1.3	−1.8	−2.2	0.0
Manufacturing	2.6	3.2	1.7	2.3
Durables	3.6	4.1	2.2	3.5
Nondurables	1.6	2.4	1.0	0.6
Transport., publ. util.	2.1	2.8	1.0	1.6
Wholesale trade	1.4	1.1	0.9	2.0
Retail trade	−0.2	−0.0	−0.1	0.0
Finance and insur.	−2.2	−2.1	−0.5	−3.0
Other services	−0.9	−1.0	0.4	−1.3
Measurable sector	2.1	2.6	0.7	2.2
Unmeasurable sector	−0.7	−0.7	−0.3	−0.8

percentage point overall; but correcting for a compatibility problem during the 1973–1979 period increases measured TFP growth by about half a percentage point. Moreover, the measurable sector appears to be more affected by this correction than the unmeasurable sector.

An increase in the growth rate of investment-specific technological change is not the only possible trigger of a learning-induced slowdown in productivity growth. Our argument can be also applied to an investment boom whether or not this boom is associated with technological change. For example, Young (1992) documents that Singapore experienced considerable growth in investment rates and capital accumulation, but no recorded TFP growth.[28] From our perspective, rapidly increasing investment would *induce* a decline in recorded TFP growth, and given the magnitude of the increase in capital accumulation in Singapore, zero measured TFP growth does not seem surprising. In contrast, however, note that an increase in the capital stock which results from an increase in the rate of technological change would be more severe, since it would also involve compatibility problems between new and old technologies.[29]

28. Investment rates increased from 9% of GDP in 1960 to 43% of GDP in 1984.

29. As a gauge on the role of the assumption that labor is freely mobile across vintages, we also considered the vintage formulation used in Cooley, Greenwood, and Yorukoglu (1995). With the same setup as above, suppose that once capital is allocated to a plant,

3.2 SLOWDOWNS DUE TO QUALITY MISMEASUREMENTS

We have argued that for most industrialized economies the unmeasurable sector, that is, the sector where output is badly measured, is large. There is a presumption that we underestimate output growth in this sector, because it is more difficult to capture quality changes of goods in this sector. We now describe a simple model in which output does have a quality and a quantity component, and we make the extreme assumption that we can only measure the quantity component of output. We then study the quantitative implications for productivity measurement when the rate of investment-specific technological growth changes in the model economy.[30]

We formalize the distinction between the quantity and quality components of output by identifying each component with a separate production process. We thus postulate that output y at a plant can be decomposed to read YQ, where Y is the number of goods, which is well measured, and Q is a one-dimensional quality index per good, which is not measured at all. Capital and labor are inputs to the production of quantity and quality

$$Y = \left(\gamma k_Y^{\alpha_Y} l_Y^{1-\alpha_Y}\right)^{\beta} \quad \text{and} \quad Q = \left(\gamma k_Q^{\alpha_Q} l_Q^{1-\alpha_Q}\right)^{1-\beta},$$

and the capital intensity (share) in the production of quantity and quality may differ. The parameter β represents the relative quantity content of the output. Note that the production technology for output measured in quality units has constant returns to scale in the capital and labor in-

it remains there until the plant shuts down, at which point capital depreciates completely. Also suppose that each plant has a fixed labor requirement of one. All labor is paid the same wage rate, and capital is allocated to a plant until the present value of the sequence of marginal products of capital equals the current cost of investment. There will be a point at which a plant shuts down, since the marginal product of labor in a given plant will grow at a slower rate than the wage rate. In order to simplify the characterization of the optimal investment decision, assume also that the interest rate is constant. A balanced-growth path for this economy can be summarized by (1) a constant growth rate for the wage rate, and (2) a fixed life span of plants. The wage rate at a point in time has to be such that the oldest vintage finds it profitable not to close down. Finally, the amounts of investment and labor attracted to new plants have to be such that present-value profits are zero when the total amount of labor hired for new plants equals the number of laid-off workers from old plants shutting down. Our analysis of this model framework leads to results which are very similar to those obtained in the model with aggregation. The qualitative features of the model are the same, and the quantitative results from the same set of parameters for learning are also very similar. We detected the largest discrepancy in the average age of firms, but this difference was not large enough to generate significantly different aggregate output paths.

30. For a related analysis of unobserved quality in an endogenous growth context see Howitt (1995).

puts.[31] We assume that in the national income accounts of this model economy only the quantity component of the consumption aggregate is measured, but that investment goods are well measured.

We will model structural change in two complementary ways, and each way has different implications for quality mismeasurement. First, we will consider the kind of experiment that we studied in the previous section: at a point in time, the growth rate of investment-specific technological change increases permanently. Second, we consider an experiment where this technological shift also leads to a relative shift toward quality production. We now develop the specific time-series implications of the experiments, first by considering the mismeasurement on the plant (vintage) level, and then on the aggregate level.

3.2.1 Quality Mismeasurement on the Plant Level Consider an isolated production facility, and note that the optimal allocation of a given amount of capital k and labor l across quantity and quality production has to satisfy

$$\frac{k_Y}{k_Q} = \frac{\beta}{1-\beta}\frac{\alpha_Y}{\alpha_Q} \quad \text{and} \quad \frac{l_Q}{l_Y} = \frac{\beta}{1-\beta}\frac{1-\alpha_Y}{1-\alpha_Q},$$

where $k_Y + k_Q = k$ and $l_Y + l_Q = l$. This implies, after some manipulations, that total output can be written

$$y = A\gamma k^{\hat{\alpha}} l^{1-\hat{\alpha}},$$

where $\hat{\alpha} \equiv \beta\alpha_Y + (1-\beta)\alpha_Q$ and $A \equiv A_Y A_Q$, with

$$A_Y \equiv \left[\left(\frac{\beta\alpha_Y}{\hat{\alpha}}\right)^{\alpha_Y}\left(\frac{\beta(1-\alpha_Y)}{(1-\hat{\alpha})}\right)^{1-\alpha_Y}\right]^{\beta}$$

and

$$A_Q \equiv \left[\left(\frac{(1-\beta)\alpha_Q}{\hat{\alpha}}\right)^{\alpha_Q}\left(\frac{(1-\beta)(1-\alpha_Q)}{(1-\hat{\alpha})}\right)^{1-\alpha_Q}\right]^{1-\beta}.$$

31. The constant-returns-to-scale property is assumed for convenience: it greatly simplifies decentralization of the model, and it allows the identification of relative prices with marginal products. However, it does carry some features which for several examples may appear unrealistic. First, the production of quality can often be thought of as a process where resources are devoted once and for all to develop a new product which will be available forever. Second, by implication of our formulation, if no effort (input) is devoted to quality production, then nothing is produced.

Furthermore, the subcomponents of output satisfy

$$Y = A_Y\,(\gamma k^{\alpha_Y} l^{1-\alpha_Y})^{\beta} \quad \text{and} \quad Q = A_Q\,(\gamma k^{\alpha_Q} l^{1-\alpha_Q})^{1-\beta}.$$

Parenthetically, notice from these facts that if both quantity and quality are well measured (and there is no unobserved learning), then standard growth accounting allows growth in γ, $\Delta \log z$, to be measured perfectly.

In our economy the national income accounts use quantities only to measure total output growth:

$$\Delta \log Y_t = \beta[\Delta \log z_t + \alpha_Y \Delta \log k_t + (1 - \alpha_Y)\,\Delta \log l_t],$$

so that measured TFP growth, $\Delta \log \hat{z}$, is

$$\Delta \log \hat{z}_t = \Delta \log Y_t - \hat{\alpha}\,\Delta \log k_t - (1 - \hat{\alpha})\,\Delta \log l_t.$$

It is important to note that the total capital share, $\hat{\alpha}$, is used in this calculation. From the last two equations it follows that

$$\Delta \log \hat{z}_t = \beta\,\Delta \log z_t - \alpha_Q(1 - \beta)\,\Delta \log k_t - (1 - \alpha_Q)(1 - \beta)\,\Delta \log l_t.$$

There are several important characteristics of this equation. First, if $\beta = 1$, so that there is no quality component, then clearly z would be accurately recovered with the growth accounting procedure. Second, if $\beta < 1$, then so long as all variables grow, productivity growth is underestimated because of quality improvements: $\Delta \log \hat{z} < \Delta \log z$. Moreover, an increase in the rate of growth of capital will increase the mismeasurement. So when the growth rate of investment-specific technological change increases, there will be more than a one-to-one increase in the growth rate in the capital stock, leading to a slowdown in measured productivity. Notice that for this result to obtain qualitatively, all that is needed is that some of the capital is used in the production of quality (i.e., our assumption that the capital share is equal in the production of quantity and quality is not crucial).[32]

Third, for a given β and $\hat{\alpha}$, measured productivity growth underestimates actual productivity growth more the higher α_Q is relative to α_Y, i.e., the more capital-intensive is quality relative to quantity production. An example where α_Q would be large can be found in banking services, where the introduction of computers has added significant convenience for the customer in the ability to quickly obtain information, move funds across deposits, make payments, and so on. Yet the quantity output,

32. In contrast, Howitt (1995) assumes that quality is a function of "knowledge," and not a direct function of capital and labor.

which may be measured as the number of deposits, is probably more directly related to the number of bank employees.[33]

Fourth, notice that the time path of the measured productivity growth rate will be monotone in response to a one-time, permanent increase in the growth rate of investment-specific technological change. To see this, note that the increase in the latter will cause the capital stock to grow faster, in a monotonic way, toward its new growth rate. Since the underestimate is directly proportional to the growth rate of the capital stock, $\Delta \log \hat{z}$ will decline monotonically to its new value. As we shall see, this result will be overturned when we consider the aggregation of mismeasurements over plants/vintages.

In quantitative terms, it is possible to calculate the maximal measured (long-run) productivity slowdown due to our hypothesis of an increase in the growth rate of capital as follows. First, without restrictions on $\hat{\alpha}$ or β, the largest possible slowdown would obviously equal the increase in the growth rate of capital (setting $\beta = 0$ and $\alpha_Q = 1$). However, restricting this estimate by the observed total capital share of $\hat{\alpha}$, the maximum bias from a one-percentage-point increase in the capital stock growth is given by $\hat{\alpha}$ itself, which can be achieved either by setting $\beta = 0$ and $\alpha_Q = \hat{\alpha}$ or by setting $\alpha_Q = 1$ and $\beta = 1 - \hat{\alpha}$. As an example, a one-percentage-point increase in the growth rate of investment-specific technological change together with a capital share of 0.4 implies a $1 \times 1/(1 - 0.4)$ increase in the long-run growth rate of the capital stock, and the maximal long-run slowdown in productivity is therefore 0.67%.

3.2.2 The Aggregation of Quality Mismeasurement over Plants Now suppose, in line with the formal structure of Section 3.1.1, that aggregate output is produced in plants with different capital vintages. Then we know from the plant level analysis of quality mismeasurement above that correctly measured output at time t in a τ-year-old plant satisfies

$$y_{t,\tau} = A\gamma_t k_{t,\tau}^{\hat{\alpha}} l_{t,\tau}^{1-\hat{\alpha}},$$

but that measured output at this plant, $\hat{y}_{t,\tau} \equiv Y_{t,\tau}$, is given by

$$\hat{y}_{t,\tau} = A_Y\left(\gamma_t k_{t,\tau}^{\alpha_Y} l_{t,\tau}^{1-\alpha_Y}\right)^{\beta}.$$

The allocation of labor across plants here is parallel to that in Section 3.1.1, so that $l_{t,\tau} = r_{t,\tau} l_{t,1} = l_t r_{t,\tau} / \Sigma_{s=1}^{\infty} r_{t,s}$. It follows that aggregate output

33. In the U.S. national income accounts, measured bank sector output is mainly based on employment data, and not on deposit counts or volumes.

satisfies the same expression as in that context. Hence, in the absence of measurement problems, the accounting for productivity growth would be straightforward and not give rise to errors. However, with the lack of measurement of quality on the plant level, aggregate measured output $\hat{y}_t$ equals $\Sigma_{\tau=1}^{\infty} \hat{y}_{t,\tau}$. Now note that the relative measured output at t in plants $t - \tau$ and $t - 1$, respectively, equals

$$\frac{\hat{y}_{t,\tau}}{\hat{y}_{t,1}} = \frac{A_Y \{\gamma_t [i_{t-\tau}(1 - \delta)^{\tau-1}]^{\alpha_Y} l_{t,\tau}^{1-\alpha_Y}\}^{\beta}}{A_Y \left[\gamma_t i_{t-1}^{\alpha_Y} l_{t,1}^{1-\alpha_Y} \right]^{\beta}} = r_{t,\tau}^{\beta},$$

so that

$$\hat{y}_t = y_{t,1} \sum_{\tau=1}^{\infty} r_{t,\tau}^{\beta}.$$

With the special case of equal capital intensities across quantity and quality production ($\alpha_Q = \alpha_Y$), this expression simplifies further:

$$\hat{y}_t = y_{t,1}^{\beta} \frac{A_Y}{A^{\beta}} \sum_{\tau=1}^{\infty} r_{t,\tau}^{\beta}.$$

In this case, since all variables on the right-hand side of this equation are true values (not mismeasured), we can compare directly with actual total output, which satisfies a similar equation:

$$y_t = y_{t,1} \sum_{\tau=1}^{\infty} r_{t,\tau}.$$

For illustration, let us again focus on the special case when capital intensities are the same. It is clear that since $\Sigma_{\tau=1}^{\infty} r_{t,\tau}^{\beta} \neq (\Sigma_{\tau=1}^{\infty} r_{t,\tau})^{\beta}$, we cannot exclude different transitional properties for measured output than for true output (with an = sign, the growth rates of measured output would be β times that of true output at all points in time).[34] Hence, productivity measurement as well will exhibit dynamics which are nontrivially different than those for output.

The economic interpretation of the above algebra is as follows. True output in our model has exact aggregation over vintages, so that a monotone increase in the rate of growth of the capital stock resulting from investment-specific technological change will lead to a monotone increase in the growth rate of true aggregate output—due to aggrega-

34. In the long run, of course, the distribution of labor and capital across plants does not change, and so $\Sigma_{\tau=1}^{\infty} r_{t,\tau}$ does not change, and measured output will grow exactly at β times the growth rate of true output.

tion (linearity), vintage effects are not present. Measured output, however, is not a linear function of true plant outputs, but a nonlinear one: *there are nontrivial vintage effects in the aggregation of measured output over plants.* In particular, as we shall see in our examples below, there is an initial slowdown in measured output growth relative to true output growth. Similarly, measured productivity growth will drop more during the first 10–20 years than in the long run. As the increase in the rate of growth of capital is in its initial phase, there is concentration of labor in the most recent vintages, and thus greater weight is placed where the errors are large. Only after substantial time, when the rate of growth is the same in all vintages, will the year-to-year mismeasurement be the same in all vintages, and the underestimate of productivity growth become constant.

3.2.3 Exogenous Changes in the Relative Importance of Quality The second expression of structural change that we consider is an exogenous change over time in the quality component. This is a stand-in exercise for developing a richer framework where quality is a state variable, and where a higher rate of technological change can have more fundamental effects on quality production than in our model. In a richer framework, it would for example be possible to better capture how the emergence of computers would orient the allocation of labor more toward quality in the form of the new kinds of services that computers can provide.

We do not have direct, aggregate evidence that there has been a shift toward quality. However, postwar patent data suggest that quality mismeasurement may be more important now than it was prior to the seventies. To understand why, suppose that we interpret quality mismeasurement as primarily occurring when new goods are introduced, and note that only patents for new final goods can represent a problem for growth accounting. This is true because any undermeasurement of the output of new intermediate goods will cancel with input underestimates when the same goods are used as inputs. Hence, a shift toward patents for final goods will make the quality mismeasurement more severe. On the basis of patent data from the United States and elsewhere, we split the patents into two groups: patents for new kinds of final outputs, and patents for new kinds of intermediate goods. Figure 7 shows the fraction of total patents which apply to final output goods. The figure reveals a nontrivial upward shift around the mid-seventies both for the United States data and for the international data.

Our theoretical experiment, then, is to make the parameter β decrease. This can be viewed either as an increase in the relative importance of quality in general output, or as an increase in the mismeasurement prob-

Figure 7 PATENTS FOR FINAL AND INTERMEDIATE GOODS

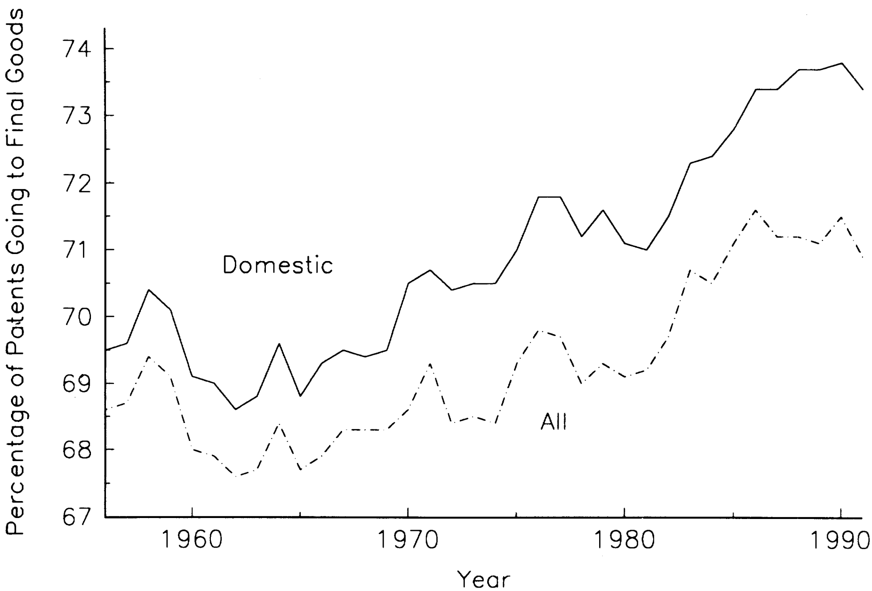

Solid curve: domestic; dashed curve: all.

lems associated with output. A change in β will have two effects. It will imply the desired move toward a larger quality component, but it will also lead to a "reduced-form" productivity increase: recall that the optimal allocation of inputs across quality and quantity components implies that a factor $\beta^{\beta}(1 - \beta)^{1-\beta}$ multiplies the aggregate production function. To isolate the first effect, we premultiply the production relations of both quality and quantity with the disembodied factor $\beta^{-\beta}(1 - \beta)^{-(1-\beta)}$. This way, changes in β will have no effect on aggregate output as correctly measured (i.e., in quality units), provided inputs are optimally allocated.

One aspect of structural change appearing in this way, as opposed to solely through an increase in the growth rate of capital, is that there may be a permanent drop in measured output growth following structural change. If there is a permanent increase in the growth rate of investment-specific technological change but no change in β, there will be a decrease in measured total-factor productivity growth, but the growth rate of output, and of labor productivity, will increase. This result is due to the simple fact that in our quality framework, where the capital share is positive (indeed the same) across the production of quantity and quality, increases in the growth rate of capital will increase the growth rates of both quality and quantity. However, when the structural change is expressed via an increase in the quality component—a decrease in β—there will be a permanent decrease in measured output and labor productivity growth rates. Our particular experiment is to assume a slow diffusion curve for β, letting it move from 0.9 to 0.8 over a period of 10 years (as the rate of investment-specific technological change increases), with the largest changes in β taking place in the middle of this period.

3.2.4 Results We perform several experiments. First, we consider the effect of a permanent increase in the rate of investment-specific technological change in an economy with unobserved quality and no learning, and where the capital intensities in quantity and quality production are equal ($\alpha_Y = \alpha_Q$). Second, we assume that, in addition to the increase in the growth rate of capital, there is a shift towards a higher quality share

Figure 8 UNOBSERVED QUALITY

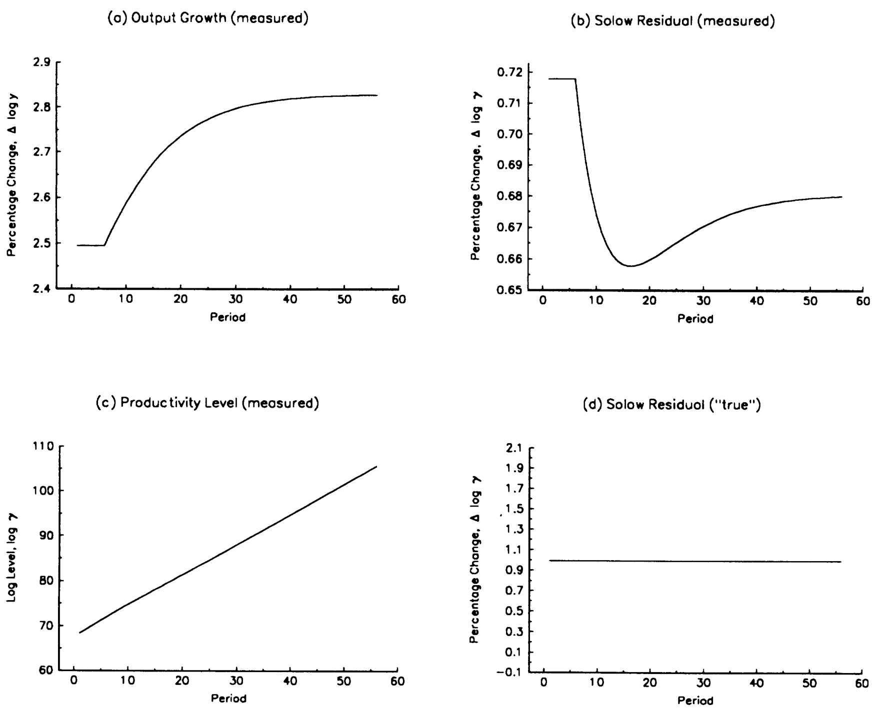

Figure 9 A SHIFT TOWARDS MORE UNOBSERVED QUALITY

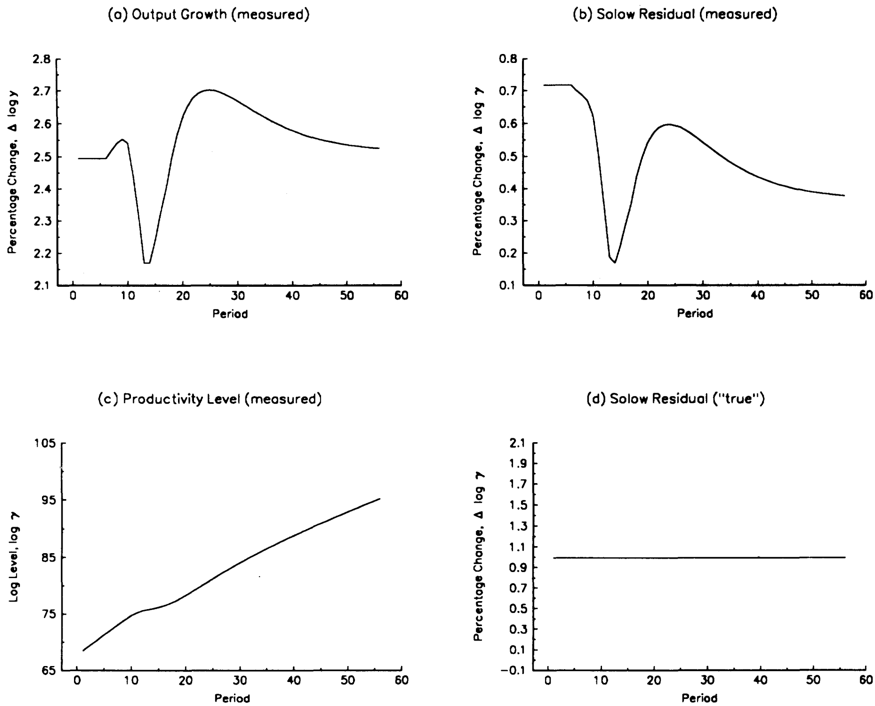

in output (a decrease in β). Third, we assume that there is learning with a compatibility problem. In the fourth and final experiment, we look at the possibility that the capital intensity in quality production is large in relative terms ($\alpha_Y < \alpha_Q$). Unless otherwise stated, the parameter values correspond to the ones in the previous section.

The results are displayed in Figures 8–11. An increase in the rate of investment-specific technological change induces more capital accumulation, and thereby an increase in the rate of growth of both observed quantity and unobserved quality. Without learning ($T_{t,1} = 1$), this results in a permanent increase in the growth rate of output and labor productivity and a permanent decline in the growth rate of measured total-factor productivity, as can be seen in Figure 8. While we observe a persistent decline in total-factor productivity growth, this effect is not quantitatively important.

The quantitative effect of a shift towards a higher quality share in output is comparable to the effect the compatibility problem has; see

Figure 10 A SHIFT TOWARDS MORE UNOBSERVED QUALITY AND LEARNING

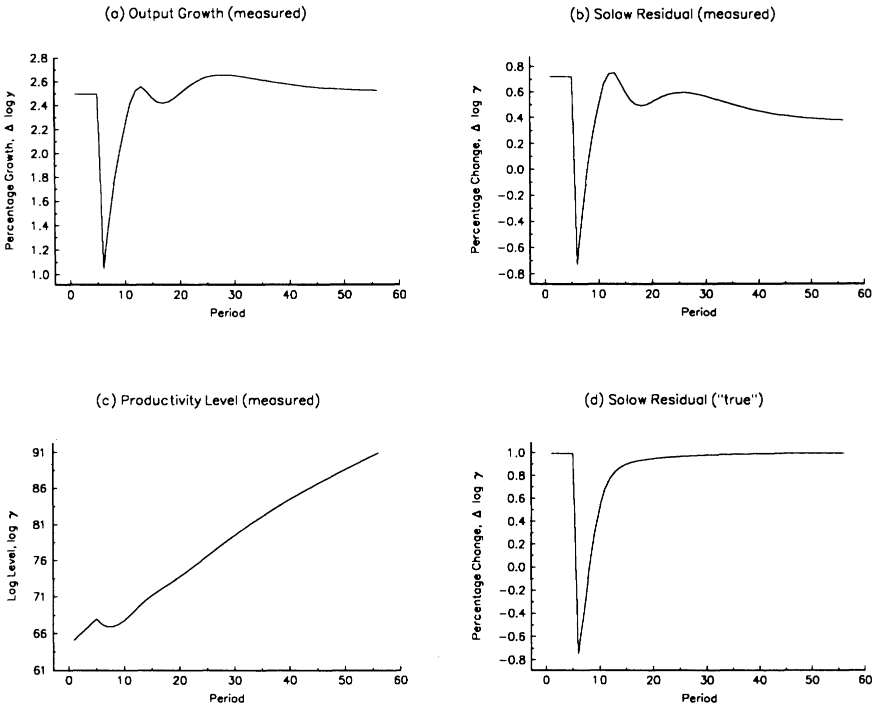

Figure 9. We observe a decline in output growth and a substantial and persistent decline in total-factor productivity growth.[35]

Adding learning with a compatibility problem does not substantially change this effect, as Figure 10 shows.[36]

Figure 11, finally, shows an example where the capital share in quality production is higher than in quantity production. We use $\alpha_Q = 1$ and $\alpha_Y = \frac{2}{9}$, so that $\hat{\alpha} = 0.3$, and we assume that β is constant, that there is no learning, and that the rates of technological change are the same as in the other figures. The graph shows that the mismeasurement increases, and that the slowdown is larger in relative terms, with otherwise similar features to the baseline case in Figure 8.

35. This persistence is to some extent due to the assumed slow diffusion of the change in β.

36. Note that the "true" Solow residual in Figure 10 is calculated by using the true output (including quality), but not adjusting the capital for learning.

Figure 11 CAPITAL IS MORE EFFICIENT IN THE PRODUCTION OF QUALITY

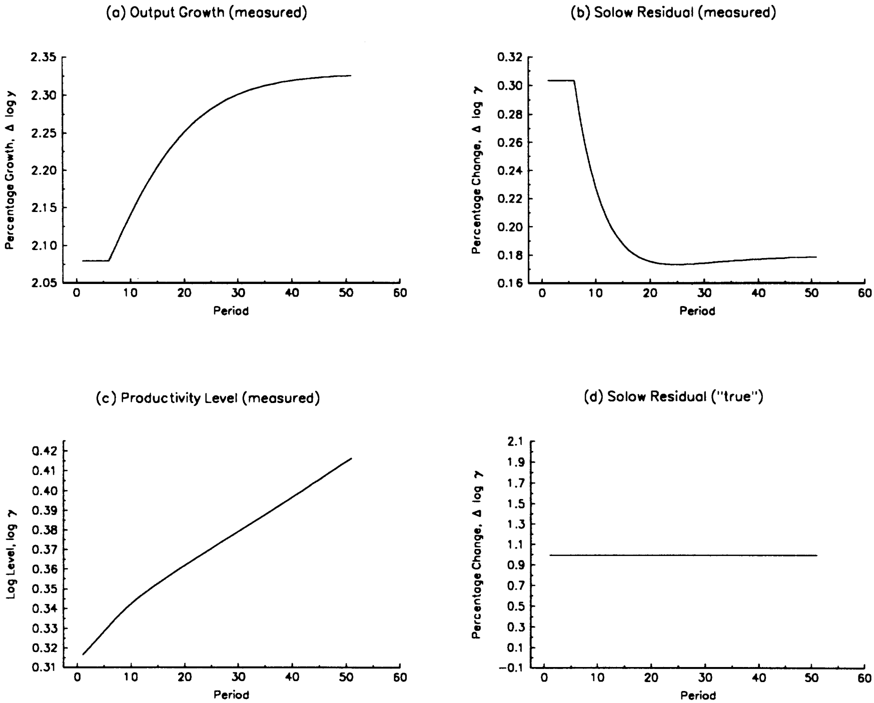

3.3 A CALIBRATED MODEL OF U.S. PRODUCTIVITY SLOWDOWN

In this section we provide a calibrated synthesis of the previous analysis using a more disaggregated model in line with the setup used at the outset of this section. Following our discussion of measurement problems in the various industries, we consider two sectors. In the measurable sector, which represents the durable- and nondurable-goods-producing industries, output is perfectly measured. In the unmeasurable sector, which represents services, production has a quality and a quantity component and measured output reflects only the quantity component. We assume that there is learning, and that the mismeasurement in the second sector is represented by the kind of mismeasurement we analyze above.

For this simulation, investment-specific technological change is not represented by a one-time regime change, but we use Gordon's relative price series for durable equipment to represent this process. Since for this experiment the rate of investment-specific technological change is not constant, we formulate the learning process as follows. The starting value for experience is $T_{t,1} = T_0 e^{-q\theta}$, where q is an unweighted 10-year moving average of past rates of investment-specific technological progress. The moving average was used in order to capture the fact that compatibility has to do with the existing skill of workers, and because workers allocated to the newest technology do not all necessarily have the most recent technological experience. For specific parameter values we chose T_0 and θ so that the starting value for the original balanced growth path is 0.8 and the starting value for the new balanced path is 0.6, the latter reflecting poorer compatibility as the improvements in capital occur at a higher rate.

The capital accumulation process in this economy corresponds to the one in the United States economy. For each sector the ratio of investment to GDP in each period equals the corresponding investment rate in the United States economy, where the measurable and unmeasurable sector are as defined in Section 2. We also equate labor income shares in the two sectors with the corresponding average labor income shares in the

Figure 12 LEARNING IN A CALIBRATED TWO-SECTOR ECONOMY

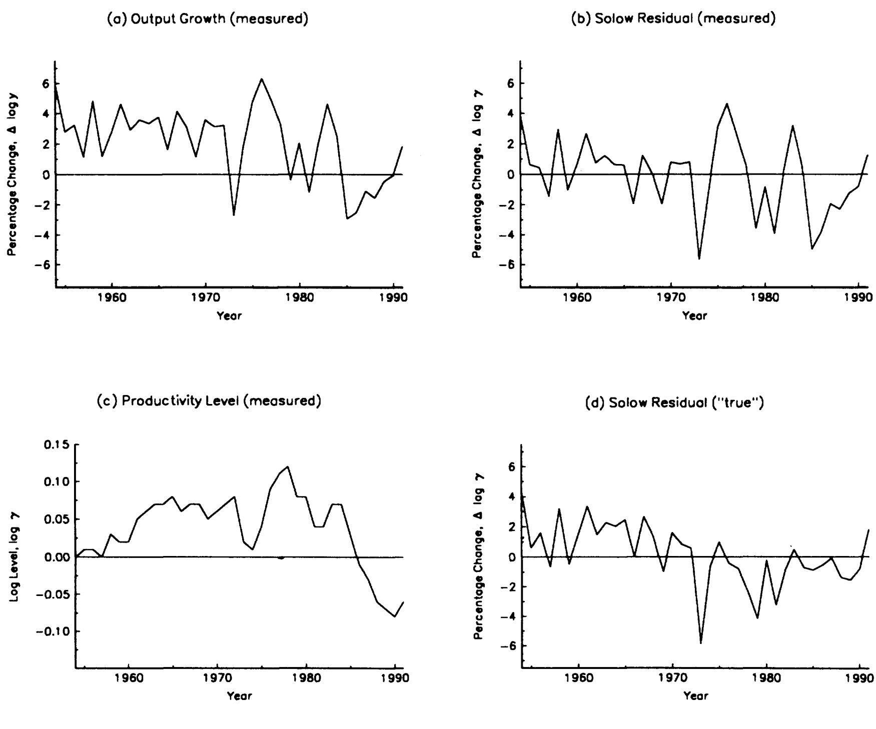

Figure 13 LEARNING AND QUALITY CHANGE IN A CALIBRATED TWO-SECTOR ECONOMY

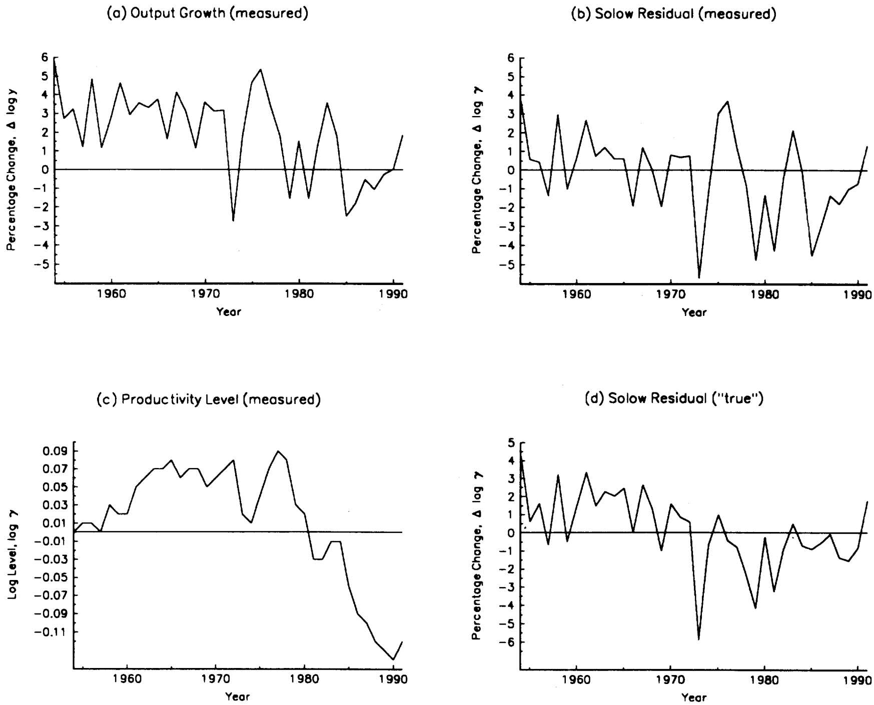

United States economy. We assume that the rate of disembodied technological change is the same in the measurable and the unmeasurable sector. The process for disembodied technical change corresponds to the neutral technical change we have calculated for the United States economy in Section 2.2.

Figure 12 displays growth rates for measured output and for measured and true disembodied technical change. Comparing parts (a) and (b), we see that the behavior of measured productivity growth is qualitatively similar to that of true productivity growth, but that it underestimates true productivity growth on average. This is reflected in part (c), which shows that the measured productivity never attains the same level as the true productivity (see Figure 3). The decline in measured productivity after 1975 is much more dramatic than the decline in true productivity. As a negative result with respect to our model's ability to replicate a productivity slowdown, we have to note that it predicts a very strong but temporary recovery immediately after 1973.

We next assume that there is also a shift towards quality that takes place in 1970. The results from this simulation are displayed in Figure 13. This feature tends to amplify the slowdown in the early 1970s, but does not have any other noticeable effects.

4. *Conclusion*

We have reviewed the productivity statistics and discussed measurement issues as well as briefly reviewed a number of suggested explanations for the productivity slowdown. As a comment on the statistics and on the literature, we then suggested a complementary candidate explanation for the slowdown: an increase in the rate of investment-specific, or capital-embodied, technological change.

The argument was presented in two steps. We first provided some evidence that there indeed has been an acceleration in technological change, and that it started around the mid-seventies, i.e., at the same time as the productivity slowdown began. We do not view this evidence as more than suggestive, but we do think there is reason to take it seriously. Second, we used a set of simple vintage capital models to discuss how an increase in the rate of capital-embodied technological change could lead to *decreases* in measured productivity growth. The channels we focused on were learning (about the new, more advanced capital) and quality mismeasurement of final outputs. We believe these two channels are real phenomena, and our task was to use "reasonable" parameterizations to get an idea of the potential quantitative importance they might have. Some of the parameters needed in these exercises are inherently difficult to assign values to. For example, we adopt learning parameters from the applied learning literature, which focuses on much more concrete tasks than those necessitated by the adoption of much of the new equipment we have in mind. Similarly, the relative importance of unmeasured quality in final output and the relative importance of the new capital in producing this quality are also parameterized on the basis of little quantitative guidance. As a result, we do not view our quantitative findings as conclusive.

We found the magnitude of a slowdown due to learning alone to be relatively small, unless the compatibility problems between the new and the old capital increase with the higher rate of technological change. However, for the recent increases in technological change, compatibility problems probably were quite important. We found the learning-induced slowdown to be relatively short-lived, about as short-lived as the slowdown recorded in the measurable sectors of the economy. Finally, we pointed out that increases in the investment rates also can produce slowdowns due to learning, even in the absence of technological change, but that such slowdowns are less severe in relative terms, since they are not associated with increased compatibility problems.

The slowdown due to capital-induced quality mismeasurements also seemed to have some potential, both in terms of magnitude and duration. We also found that the quality mismeasurement story produces a temporary large drop in measured productivity, something which came about from nontrivial vintage effects. Quantitatively, we found that the effects of quality change are much more important when there is a structural change from quantity to quality production. Although we present some evidence along these lines and we do believe that the new capital introduced since the seventies (especially information capital) has made quality a more important part of output in relative terms, it is extremely difficult to make quantitative statements about the extent of this phenomenon. In sum, our experiments indicate some potential, but the low confidence that we have in some of our key parameter values prevents us from claiming either success or failure.

The shortcuts taken in the model formulation and the empirical implementation of our theory are many, and we abstain from making a wish list.[37] We do believe it is important to look more carefully at the role capital plays in improving the quality of new goods and services, and it is clearly necessary to go beyond the simple formalization we employ in the paper.

Let us also make some final remarks. First, the hypothesis of an increase in the rate of investment-specific technological change likely has implications beyond those we study here, and a look at these would give more information about the validity of the hypothesis. One such implication has already been explored in parallel work: the implication for wage inequality. For suppose, as has been documented in a number of classic studies [see, e.g., Griliches (1969) and, more recently, Flug and Hercowitz (1995), who use international data], that capital is more complementary with skilled than with unskilled labor. Then as the new capital becomes available at a higher speed, the wages of skilled agents will tend to rise relative to those of unskilled agents. This is indeed what we have observed; there was a sharp increase in wage inequality in the late seventies in the United States, and the wage gap has increased since then.[38] Krusell et al. (1995) investigate this hypothesis in detail.

Second, to the extent it is taken seriously, the view that learning and mismeasurements caused the measured productivity slowdown represents an optimistic assessment of the past and current state of affairs.

37. Besides, we prefer to stay close in spirit to the approach advocated in Romer (1992).
38. Similarly, in Europe the wage gap did not increase much, but the unemployment rate rose sharply, and did not come back down.

First, if the slowdown is due to learning, it is normal under the particular circumstances, and the implied expectation is that growth and productivity will go back up. Second, if the problem is mismeasurement of quality, then there is no need to worry either; measured decreases in the growth rate of productivity are simply misleading in this case. Notwithstanding the logic of this position, there would still be a need to deal with the widespread perception that real wages have gone down. At least for large fractions of the population, in particular the less fortunate ones, real wage growth has been dismal for a long period. Again, mismeasurement could turn this fact to a less worrisome one: it is the overestimated growth rates of price increases that cause the slowdown in real wage growth, and properly adjusted, real wage growth has not been as catastrophic. However, as suggested above, it may also be that the real wage decline has not just been a statistical artifact, and that it reflects the kind of distributional change which may follow from an increase in the rate of investment-specific technological change.

Appendix. Data Sources

Industry output is value-added or gross-product-originating (GPO) in constant dollars from the BEA. In 1991 the BEA revised industry GPO data substantially, and it is now publishing constant-1987-dollar industry GPO data starting in 1977. To obtain industry output series for the period 1954–1993, we have linked the constant-1982-dollar prerevision industry GPO series from 1954 to 1976 with the current revised series in 1977. Industry employment is the number of full-time equivalent employees. Equipment and structures capital stocks are constructed from industry investment data from the BEA, assuming constant geometric depreciation. The annual depreciation rates are 12.4% for equipment and 5.6% for structures. For structures we use constant-1987-dollar investment data. For equipment we use current-dollar investment data and deflate the series with the PDE price index or Gordon's equipment price index. When we correct for the effects of learning, we construct a capital stock index for equipment as described in Section 3.1.1, assuming a learning rate of $\lambda = 0.7$ and initial experience of $T_0 = 0.8$ before 1973 and $T_0 = 0.6$ after 1973.

With one exception, we assume that the total capital and labor shares are constant throughout the period. We calculate average income shares from 1947 to 1987 based on updated data from Jorgenson, Gollop, and Fraumeni (1987). Moreover, we assign shares to equipment and structures according to Greenwood, Hercowitz, and Krusell (1995), where these shares are selected to match long-run return equalization between the two types of capital. As a result, equipment is assigned 57% of the total capital share and structures 43%. This is close to findings based on long-run cost data in Duménil and Lévy (1990). We apply the same capital-share division in all sectors. The capital share for the whole private sector implied by these sectoral shares is 33%. Jorgenson's capital-share numbers are somewhat high in comparison with what others have used; for example, Gordon (1990) uses 25%.

We do not assume that the capital and labor income shares are constant when we construct the productivity of the PDE sector relative to that of the remainder of the economy in Section 2.4. We use the income components of industry GPO from the BEA to define capital and labor income shares. Not all income components can be unambiguously assigned to capital or labor income. We assume that from all sectoral income components, labor compensation can be unambiguously allocated to labor income, and interest payments, corporate profits before tax, inventory valuation adjustment, rental income, and capital consumption allowance can be unambiguously allocated to capital income. We calculate the labor income share in the sum of these income components, and assume that the labor income share in the ambiguous income components is the same as the income share calculated above. Again we assume that equipment capital receives a constant fraction of capital income, namely 0.57. To calculate the labor income share of the investment sector we identify the PDE-producing sector with the durable-goods manufacturing sector. The remaining industries are identified with the consumption-goods-producing sector. We use a three-period moving average of labor income shares. We also allocate the stock of structures in the durable-goods manufacturing sector to the PDE-producing sector and the remainder to the consumption-goods and services sector.

REFERENCES

Alchian, A. (1963). Reliability of progress curves in air-frame production. *Econometrica* 31:679–693.

Allen, S. (1985). Why construction industry productivity is declining. *Review of Economics and Statistics* 67(4):661–669.

Bahk, B.-H., and M. Gort (1993). Decomposing learning by doing in new plants. *Journal of Political Economy* 101(4):561–583.

Baily, M. N. (1981a). Productivity and the services of capital and labor. *Brookings Papers on Economic Activity*, pp. 1–50.

Baily, M. N. (1981b). The productivity growth slowdown and capital accumulation. *American Economic Review Papers and Proceedings* 71(2):326–331.

Baily, M. N., and A. Chakrabarti. (1988). *Innovation and the Productivity Crisis.* Washington: The Brookings Institution.

Baily, M. N., and R. Gordon. (1988). The productivity slowdown, measurement issues, and the explosion of computer power. *Brookings Papers on Economic Activity*, pp. 347–420.

Baumol, W., and E. Wolff (1979). On the theory of productivity and unbalanced growth. C.V. Starr Center for Applied Economics. New York University.

Boskin, M., E. Dulberger, Z. Griliches, R. Gordon, and D. Jorgenson. (1995). Toward a more accurate measure of the cost of living. Interim Report to the Senate Finance Committee from the Advisory Commission to Study the Consumer Price Index.

Bruno, M. (1981). Raw materials, profits, and the productivity slowdown. *Quarterly Journal of Economics* 99(1):1–29.

Bruno, M., and J. Sachs. (1982). Input price shocks and the slowdown in economic growth: The case of U.K. manufacturing. *Review of Economic Studies* 49:679–705.

Christiansen, G., and R. Haveman. (1981). Public regulations and the slowdown in productivity growth. *American Economic Review Papers and Proceedings* 71(2):320–325.

Cooley, T., J. Greenwood, and M. Yorukoglu. (1995). The replacement problem. Institute for Empirical Macroeconomics. Federal Reserve Bank of Minneapolis.

Cullison, W. (1989). The U.S. productivity slowdown: What the experts say. *Federal Reserve Bank of Richmond Economic Review* 75(4, July/August):10–21.

Darby, M. (1984). The U.S. productivity slowdown: A case of statistical myopia. *American Economic Review* 74:301–322.

Dean, E., K. Kunze, and L. Rosenblum. (1988). Productivity change and the measurement of heterogeneous labor inputs. Mimeo.

De Leeuw, F., M. Mohr, and R. Parker. (1991). Gross product by industry, 1977–88: A progress report on improving the estimates. *Survey of Current Business* 71:23–37.

Denison, E. (1985). *Trends in American Economic Growth, 1929–1982*. Washington: The Brookings Institution.

Dertouzos, M., R. Lester, and R. Solow. (1990). *Regaining the Productive Edge*. Cambridge, MA: Harper-Perennial, The MIT Commission on Industrial Productivity.

Duménil, G., and D. Lévy. (1990). Continuity and ruptures in the process of technological change. CEPREMAP. Mimeo.

Edquist, C., and S. Jacobsson. (1988). *Flexible Automation—The Global Diffusion of New Technology in the Engineering Industry*. Oxford: Basil Blackwell.

Flug, K., and Z. Hercowitz. (1995). Some international evidence on equipment-skill complementarity. Department of Economics, Tel Aviv University.

Gjerding, A. (1991). Produktivitetsparadokset og den Nye Teknologi. *Nationalokonomisk Tidskrift* 129:54–66.

Gordon, D. M. (1981). Capital–labor conflict and the productivity slowdown. *American Economic Review Papers and Proceedings* 71(2):30–35.

Gordon, R. (1990). *The Measurement of Durable Goods Prices*. Chicago: National Bureau of Economic Research, University of Chicago Press.

Greenwood, J., Z. Hercowitz, and P. Krusell. (1996). Long-run implications of investment-specific technological change. Rochester Center for Economic Research. University of Rochester. Forthcoming in *American Economic Review*.

Greenwood, J., and M. Yorukoglu. (1996). 1974. Paper presented at the Carnegie-Rochester Conference on Public Policy, Spring 1996.

Griliches, Z. (1969). Capital–skill complementarity. *Review of Economics and Statistics* 51:465–468.
Griliches, Z. (1988). Productivity puzzles and R&D: Another nonexplanation. *Journal of Economic Perspectives* 2(4):9–21.
Griliches, Z. (1994). Productivity, R&D, and the data constraint. *American Economic Review* 84:1–23.
Gullikson, W. (1995). Measurement of productivity growth in U.S. manufacturing. *Monthly Labor Review,* July, pp. 13–20.
Hayes, R., and W. Abernathy. (1980). Managing our way to economic decline. *Harvard Business Review,* 58:67–77.
Howitt, P. (1995). On some problems in measuring knowledge based growth. Paper presented at the conference on "Implications of knowledge-based growth for micro-economic policies", Ottawa 1995.
Hulten, C., J. Robertson, and F. Wykoff. (1989). Energy, obsolescence, and the productivity slowdown. In D. Jorgenson and R. Landau (eds.). *Technology and Capital Formation,* Cambridge, MA: The MIT Press, pp. 225–258.
Hulten, C. R. (1992). Growth accounting when technical change is embodied in capital. *American Economic Review* 82:964–980.
Jorgenson, D. W., F. M. Gollop, and B. M. Fraumeni. (1987). *Productivity Growth and U.S. Economic Growth.* Harvard University Press.
Jovanovic, B., and Y. Nyarko. (1995). A Bayesian learning model fitted to a variety of empirical learning curves. *Brookings Papers on Economic Activity—Microeconomics,* 247–305.
Kendrick, J. (1987). Service sector productivity. *Business Economics* 22(2):18–24.
Kendrick, J. (1990). International comparisons of productivity trends and levels. *Atlantic Economic Journal* 18(3):42–54.
Krusell, P., L. Ohanian, J.-V. Ríos-Rull, and G. Violante. (1995). Capital–skill complementarity and inequality. University of Rochester. Mimeo.
Lindbeck, A., and D. J. Snower. (1995). Restructuring production and work. University of Stockholm. Institute for International Economics.
Lysko, W. (1995). Manufacturing multifactor productivity in three countries. *Monthly Labor Review,* July, pp. 39–55.
McHugh, R., and J. Lane. (1987). The role of embodied technical change in the decline of labor productivity. *Southern Economic Journal* 53(4):915–924.
Naples, M. (1981). Industrial conflict and its implications for productivity growth. *American Economic Review Papers and Proceedings* 71(2):36–41.
Naples, M. (1988). Industrial conflict, the quality of worklife, and the productivity slowdown in U.S. manufacturing. *Eastern Economic Journal* 14(2):157–166.
Nordhaus, W. (1982). Economic policy in the face of declining productivity growth. *European Economic Review* 18(2):131–157.
Norsworthy, J., M. Harper, and K. Kunze. (1979). The slowdown in productivity growth: Analysis of some contributing factors. *Brookings Papers on Economic Activity,* pp. 387–421.
Olson, M. (1988). The productivity slowdown, the oil shocks, and the real cycle. *Journal of Economic Perspectives* 2(4):43–69.
Rapping, L. (1965). Learning and World War II production functions. *Review of Economics and Statistics* 47:81–86.
Romer, P. M. (1992). Crazy explanations for the productivity slowdown. *NBER Macroeconomics Annual,* pp. 163–210.

Shapiro, M. D., and D. Wilcox. (1996). Mismeasurement in the consumer price index: An evaluation. *NBER Macroeconomics Annual,* this issue.
Solow, R. M. (1957). Technical change and the aggregate production function. *Review of Economics and Statistics* 39:312–320.
Solow, R. M. (1959). Investment and technical progress. In K. Arrow, S. Karlin, and P. Suppes (eds.). *Mathematical Methods in the Social Sciences,* Stanford, CA: Stanford University Press.
Summers, L. (1982). Comments on "Economic policy in the face of declining productivity growth" by Nordhaus. *European Economic Review* 18(2):163–169.
Thurow, L. (1981). Solving the productivity problem. In L. Thurow, A. Packer, and H. Samuels (eds.). *Strengthening the Economy: Studies in Productivity,* Washington: Center for Democratic Policy, pp. 9–19.
Thurow, L. C. (1987). Can America compete in the world economy? In Y. Shetty, and V. M. Buehler. Quality, Productivity, and Innovation: Strategies for Gaining Competitive Advantage. New York: Elsevier Science.
Yelle, L. (1979). The learning curve: Historical review and comprehensive survey. *Decision Sciences* 10:302–328.
Yorukoglu, M. (1995). The information technology paradox. Mimeo.
Young, A. (1992). A tale of two cities: Factor accumulation and technical change in Hong Kong and Singapore. *NBER Macroeconomics Annual,* pp. 13–54.

Comment

ROBERT J. GORDON
Northwestern University and National Bureau of Economic Research

1. Documenting the Slowdown

The paper by Hornstein and Krusell (hereafter HK) documents the productivity-growth slowdown, treats some of the related measurement issues, and introduces a theoretical explanation of the slowdown based on two apparently perverse channels through which increases in the rate of investment-specific technological change can lead to decreases in measured productivity growth. This comment qualifies HK's discussion of the facts and measurement issues and questions the plausibility of the theoretical explanation.

The authors document the productivity-growth slowdown, which began in the late 1960s and exhibits a sharp transition to slower growth around 1972–1973. Unfortunately, HK's tables are based on official data that were made obsolete by important revisions released in early 1996. The revisions shift all official U.S. government output and productivity data from a fixed 1987 base year (which has the effect of underweighting computers prior to 1987 and overweighting computers after 1987) to chain weights that more accurately reflect the shifting relative prices and quantities of computers. The effect of the revisions is to decrease substantially the growth rate of labor productivity and TFP after 1987 and increase both rates prior to 1987.

The revised data, presented in my Table 1, provide a quite different picture of productivity acceleration and deceleration. TFP growth in the aggregate (private business) economy exhibits a marked deceleration between 1972–1979 and 1979–1994 in the revised data, in contrast to the acceleration for the aggregate economy reported by HK from unrevised data in the top line of their Table 2. This further deceleration is important for distinguishing among hypotheses to explain the slowdown. Factors specific to the 1970s (e.g., oil prices) now lack conviction as explanations, whereas factors that intensified in the 1980s and 1990s (possibly measurement problems) may become more convincing.

The revisions change the growth rate of TFP in the manufacturing sector relative to the total private economy (and thus by inference to the "residual" nonfarm nonmanufacturing sector, i.e., (NFNM). In HK's Table 2, TFP growth in manufacturing for the most recent period, 1979–1993, exceeded TFP growth in the private economy by 0.9 percentage points (1.6% vs. 0.7%), about the same as in the revised data (1.19% vs. 0.23%). But during the 1970s the relationship is quite different. Whereas in HK's Table 2 TFP growth for manufacturing performed better than in the private economy (−0.2% vs. −0.6%), in my Table 1 growth in manufacturing performed substantially worse (0.33% vs. 0.85%). The sharp downward revision of TFP growth, which is concentrated outside of

Table 1 DIMENSIONS OF THE U.S. PRODUCTIVITY SLOWDOWN, 1950–1972 VS. 1972–1994

Concept and sector	*1950–1972*	*1972–1979*	*1979–1994*	*1972–1994*	*1972–1994 minus 1950–1972*
Output per hour					
Private business sector	3.12	1.50	1.22	1.31	−1.81
Nonfarm business	2.65	1.30	1.06	1.13	−1.52
Manufacturing	2.62	2.17	2.62	2.47	−0.15
Nonfarm nonmanufacturing	2.66	0.95	0.57	0.67	−1.99
Total Factor Productivity (TFP)					
Private business sector	1.93	0.85	0.23	0.43	−1.50
Nonfarm business	1.64	0.67	0.05	0.25	−1.39
Manufacturing	1.54	0.33	1.19	0.90	−0.64
Nonfarm nonmanufacturing	1.69	0.81	−0.31	0.03	−1.66

Sources: Bureau of Labor Statistics and author's calculations.

manufacturing, makes obsolete the authors' observation that the United States "is among a minority [of countries] showing partial improvement in the 1980s."

What do these data tell us about the contribution of capital to the productivity slowdown? Simple algebra shows that the growth of capital per hour can be computed from Table 1 as the difference between the growth rate of output per hour and the growth rate of TFP, divided by the income share of capital. This calculation reveals the striking fact that growth in capital per hour actually *accelerated* in manufacturing after 1972, indicating that capital made a *negative* contribution to the slowdown in labor productivity growth in manufacturing. In NFNM both capital per hour and TFP share the blame for the slowdown, although TFP accounts for 85% of it.

2. *Measurement Issues*

While measurement problems have been much discussed, until recently they did not seem to be a plausible explanation for the TFP growth slowdown, simply because measurement problems "have always been with us." Even if we could come up with reasons to believe that TFP growth had been understated by, say, 1.0 percentage point per annum since 1972, the same measurement issues would make the understatement for the pre-1972 period just as great. The magnitude of the slowdown would remain unchanged (although the *percentage* decline in the growth rate would be substantially reduced).

Two factors have emerged recently that provide a contribution to explaining the slowdown. First, as Griliches (1994)[1] has pointed out and HK document in their Tables 4 and 5, the share of the economy where output is "unmeasurable" (i.e., difficult to measure) has increased substantially, from 48% in 1954 to 66% in 1993. This compositional effect can explain 0.46 percentage points of the productivity-growth slowdown between 1954–1973 and 1979–1993, according to HK's Table 4.[2] This is a striking finding, since the compositional shift seems to explain 46/60, or 77% of the total slowdown over those two intervals. This is a much larger share of the slowdown than has been explained by any other factor proposed in the literature.

Unfortunately, the slowdown to be explained is much greater in the revised data, as shown in my Table 1. Between 1950–1972 and 1979–

1. References in this comment are to the list in HK.
2. Comparing 1979–1993 with 1954–1973, the aggregate slowdown in Table 4 is 0.6, and the average slowdown in the two sectors (taking an average of the 1954 and 1993 weights from Table 5) is 0.14, so that the compositional effect explains the difference (0.46).

1994, the slowdown for the private business sector in my Table 1 is 1.70 points per annum, not 0.60 as in HK's Tables 2 and 4 for the comparable periods. While we do not have revised data that would allow us to compute the compositional effect, it will surely amount to one-third or less of the total slowdown that requires an explanation.

Another measurement contribution is the so-called "formula bias" in the CPI. As discussed in the interim report of the CPI commission (Boskin et al. 1995), this amounts to 0.5% per annum and applies to the period from 1978 to 1995. Because it does not apply before 1978, it helps to explain the productivity-growth slowdown in the post-1978 period relative to pre-1978. Multiplying by the share of consumption (net of imputed rent on owner-occupied dwellings) in total private business output (about 78%), the formula bias could account for a 0.4% per-annum downward bias in the growth rate of private output and in private-sector TFP.

Combined with the compositional effect, these two factors could explain as much as 0.9 percentage points of the slowdown, a very substantial amount. The two factors support and extend the authors' conclusion that measurement problems are large and have increased substantially during the postwar period.

HK introduce another measurement issue, the adjustment based on my price research, which they apply to the output of durable goods and the stock of durable goods. As they point out, any correction for a secular upward bias in conventional price indexes for durable goods has a two-sided effect on TFP by simultaneously boosting the output of durable goods industries and raising the growth rate of the capital of producers' durables. Their correction for this upward bias, based on numbers from my book and extrapolations therefrom, boosts TFP growth in the durable-goods sector by 2 percentage points per annum while cutting TFP growth in the remaining sectors by 0.3–0.4 percentage point per annum.

My price indexes for durables grow more slowly than official indexes for a host of reasons; the authors err in attributing the difference entirely to quality change and err even more by equating "quality adjustments" to the concept of "technological change embodied in new durable goods." There are two different points, both of which may be understood by assuming a society entirely devoid of technological change. First, the official price indexes may err in measuring the prices of goods of constant quality for a host of reasons, including traditional substitution bias when the relative prices of two goods change for reasons having nothing to do with technological change. Second, the official price indexes may err by neglecting quality change in a technologically stagnant society. For instance, there may be a technological frontier which allows the production of a variety of different refrigerators ranging from inexpensive and energy-inefficient to expensive and energy-efficient. A change in energy prices may cause consumers to shift from the former to the latter, and the official indexes may miss entirely or understate the extent to which the resulting price increase actually represents an improvement in quality. But because there has been no technical change by assumption, just a shift in the mix of energy-inefficient and -efficient models, any error in the official price indexes has no implication for the rate of technological change. This type of quality change is sometimes called "cost-increasing" quality change, because the energy-efficient models are more costly to produce at a given level of technology but will be voluntarily purchased if the level of energy prices is sufficiently high.

In principle, then, the differing secular growth rates of my alternative price indexes and the official price indexes for durable goods provide no evidence at all on the rate of equipment-embodied technological change. This point is of theoretical importance, but does not qualify the nature of the authors' quantitative adjustments to TFP; if my alternative price indexes are correct and the official indexes are incorrect, then official measures of real output and capital input should be adjusted by the difference in the growth rates of the two price indexes, no matter what the cause of that difference.

3. *Can a Technological Acceleration Cause a TFP Deceleration?*

The authors' basic (and perverse) hypothesis is that an acceleration in the rate of investment-specific technological change can lead to a decrease in measured TFP growth. They base their hypothesis of an acceleration in technological change on Figure 2, which plots the price of PDE relative to nondurable consumption goods and services. However, Figure 2 does not support their basic presumption of an acceleration in technical change. Clearly the evolution of that price ratio was influenced by the oil shocks of the 1970s. Taking intervals that are not influenced by the oil shocks, we can calculate that the annual rate of growth of the price ratio in Figure 2 was −3.5% from 1953–1970 and −3.4% for 1970–1993, thus providing no evidence of an acceleration of the rate of decline.

The authors develop and simulate a model to assess the effects of a technological acceleration. Leaving aside the unconvincing nature of their evidence from Figure 2 on the existence of this acceleration, none of their simulations helps us understand the productivity slowdown. Neither of the learning simulations in Figure 5 or 6 provides any explanation of a slowdown in the growth rate of the Solow residual. In both cases there is a sharp transitory drop in the residual followed by a recovery to the original level. During the recovery period the residual rises, i.e., grows faster than in the base period prior to the shock. Thus to be consistent with the author's learning simulations we should have observed a sharp deceleration of TFP growth followed by an *acceleration* to a growth rate substantially above the pre-shock (i.e., pre-1970) growth rate. No such time path has been observed for the actual behavior of U.S. TFP. Instead, as shown in my Table 1, there was a two-stage slowdown, initially during 1972–1979 and then a further slowdown in 1979–1994.

Several experiments are conducted with a related "mismeasurement" model. In all four experiments plotted in Figures 8–11 the "true" Solow residual is flat (with a temporary downward spike in Figure 10), but the incorrectly measured Solow residual exhibits a permanent decline. While the details differ across the experiments, the essential cause of the absolute decline in the residual is that the effect of technological change in increasing capital input is better measured than the resulting increase in output caused by that capital input.

To assess the plausibility of the mismeasurement results, I prefer to use a much simpler model that brings out an important flaw in HK's exercises. The growth in output (Δy) is a weighted average of growth in consumption (Δc) and investment goods (Δi). Thinking about a steady state in which we can neglect the distinction between the growth of investment and capital (Δk), we have the growth of aggregate output as

$$\Delta y = \beta\, \Delta c + (1 - \beta)\, \Delta k. \tag{1}$$

The growth in consumption goods (and services) is equal to the growth in the Solow residual in the consumption section (Δs_c) plus a weighted average of the growth in labor and capital inputs. Assuming for convenience that the growth rate of labor input is zero, the production process for consumption goods is

$$\Delta c = \Delta s_c + \alpha\, \Delta k. \tag{2}$$

Finally, the growth in investment (and in capital) is produced by the same production process, with a different Solow residual (Δs_k):

$$\Delta k = \Delta s_k + \alpha\, \Delta k = \frac{\Delta s_k}{1 - \alpha}. \tag{3}$$

Table 2 EFFECTS OF MISMEASUREMENT IN A SIMPLE MODEL[a]

	Output growth Δy	*Capital growth* Δk	*Residual growth* Δa
1. Initial situation	4.00	4.00	2.80
2. Accelerate Δs_k to 3.5			
a. True value	4.44	5.00	2.94
b. Mismeasure Δk	4.24	4.00	3.04
c. Mismeasure Δc	4.20	5.00	2.66

[a] Assumed parameters $\alpha = 0.3$, $\beta = 0.8$, "true" $\Delta s_c = \Delta s_k = 2.8$.

Since labor input is constant, the aggregate value of the Solow residual can be calculated as $\Delta a = \Delta y - \alpha \Delta k$. We may easily use this simple structure to evaluate the HK proposition that mismeasurement can cause an acceleration in capital-embodied technical progress to cause a decline in the Solow residual. We begin in line 1 of my Table 2 with an initial situation prior to the technological acceleration. The calculations take the share of capital in income to be $\alpha = 0.3$ and the share of consumption in output to be $\beta = 0.8$. The initial value of the Solow residual in both the consumption and the investment sector is assumed to be 2.8, and the growth rate of both output and capital is 4.0. Now let us raise the Solow residual in the investment sector from 2.8 to 3.5. The economy's true response is shown in line 2a, with an acceleration of capital growth from 4.0% to 5.0%, in output growth from 4.0% to 4.44% and in the aggregate Solow residual from 2.8% to 2.94% (since the share of investment in output is 0.2, the aggregate residual accelerates by 0.14, which is 0.2 times the 0.7 acceleration in the capital goods sector).

How does mismeasurement change this story? If mismeasurement causes the acceleration in capital-goods output to be missed, but consumption output (including the contribution of capital to consumption) is measured correctly, we obtain line 2b. The Solow residual accelerates, not the result we are looking for. Instead, we need the opposite type of mismeasurement, in which the output of investment goods and the input of capital are measured accurately, but the contribution of the extra capital growth to consumption growth is entirely missed. This yields line 2c, with a deceleration in the Solow residual from the initial 2.80 in line 1 to 2.66.

But here the simple example reveals the flaw in HK's exercise. In line 2c we manage to obtain a deceleration of the Solow residual, but only in a situation in which the measured growth in capital *accelerates*. And this is counterfactual, as shown in Table 1. Since the growth rate of the capital–labor ratio can be calculated as $(\Delta y - \Delta n - \Delta a)/\alpha$, it appears that growth in capital decelerated from 4.0% to 2.9% per year in the total private economy between 1950–1972 and 1972–1994, and from 3.2% to 2.1% per year in the NFNM sector. While, as pointed out above, the growth of the capital–labor ratio did accelerate in the manufacturing sector, this is of no help to the authors, since what they are trying to explain in this paper is the productivity-growth slowdown outside of manufacturing.

A further problem is that there is no connection between the theoretical exercise and the data section. Have the sectors exhibiting the largest TFP growth slowdowns been those experiencing the greatest acceleration in capital quality? In some sectors, particularly communications (not shown separately in HK's Table 2), capital quality, output growth, and the Solow residual have all accelerated. Perhaps the leading candidate for capital acceleration and a Solow residual slowdown is financial services, where there has been a massive investment in computers but where the payoff in terms of higher transactions volume and improved quality has been almost entirely missed in official output measurement.

4. *An Alternative Explanation: The Slowdown is Real*

As we have seen, there is scant if any evidence of a technological acceleration, and it is impossible to concoct a model scenario in which such an acceleration causes a TFP slowdown without also causing a counterfactual acceleration in the growth rate of measured capital. There is little if any evidence available that the measurement of consumption-goods output is much worse than capital-goods output, as is required by HK's mismeasurement hypothesis. HK's exercises are inconsistent with the basic facts of the productivity slowdown, which are that the growth rates of output per hour, of capital per hour, and of the Solow residual *all* slowed down and by about the same amount, and that the same pattern is observed across all countries. Surely mismeasurement is important whether there has been a technological acceleration or not, and the combination of a shift in output toward the poorly measured part of the economy, plus the CPI "formula bias," may explain part of the slowdown.

But much of the rest of the slowdown may be real. It is quite possible that the great invention of the last part of this century, the electronic computer, does not have the potential to achieve a massive increase in TFP as did earlier inventions. It is a stylized fact that TFP growth remained at or below 0.5% per annum during nineteenth century, accelerated to 1.5% between 1915 and 1965, and then has decelerated back to 0.5% or below since the late 1970s. The "one big wave" of American economic growth during 1915–1965 reflects the combined influence of several central inventions that, taken together, had a much more profound impact on the way the economy and society operated than has the electronic computer. These great inventions of the early twentieth century include:

the pervasive spread of electric motors and appliances into all aspects of production and consumption,
the use of the internal combustion engine in motor transport and air transport, with the derivative inventions of the suburb, interstate highway, and supermarket,
the confluence of oil refining, chemicals, and plastics and their derivatives,
the telephone and its derivatives, and
the range of entertainment and information industries, including radio, movies, television, and recorded music.

Part of the reason that electronic computers have thus far failed to produce a TFP revolution is that they still represent a very small fraction of the capital stock (especially when structures are included as capital). But there is an additional sense in which computers are not as "important" as our list of early twentieth century inventions. While retrospective exercises are inevitably subjective, it is helpful to ask oneself, did America's "true" standard of living change as much between 1955 and 1995 as it did between 1915 and 1955, or between 1875 and 1915? There are plenty of television reruns and magazine advertisements that allow us to relive life in 1955, and my feeling is that it does not differ as radically from our present conditions as life in 1915 differed from that in 1955.

In many senses we are sliding down a curve of diminishing returns: the transition from LPs to CDs is not as big a deal as going from nothing to records, just as the transition from movies in the theater to movies at home on the VCR is not as big a deal as going from nothing to movies. Word processing on a Pentium computer with Windows 95 compared to a decade-ago IBM 286 represents a smaller transition than the inventions of the personal computer or the invention of the typewriter. As amazing as it may seem, we somehow managed to win World War II when only typewriters, not computers, were available to keep track of 12 million people in the armed forces. While my interpretation may seem gloomy, it seems to be supported by the marked slowdown in TFP growth in recent years in Japan, Germany, and other countries, as they converge to the American technological frontier.

Comment

VALERIE A. RAMEY
University of California—San Diego

1. *Introduction*

This paper offers a provocative new explanation for the measured productivity slowdown that began in the early 1970s. In contrast to other explanations, which emphasize *negative* economic shocks, this paper presents the case for a *positive* economic shock as the source of the productivity slowdown. The idea is quite radical: *increases* in the rate of investment-specific technological progress can manifest themselves as prolonged periods of *diminished* economic performance. To derive this result, Hornstein and Krusell emphasize two aspects of investment-specific technological progress. First, adoption of new technology requires time to learn, resulting in temporarily lower productivity. Second, technological change increases the quality of goods in a way that is very difficult to measure. The result is measurements that understate true productivity growth.

Hornstein and Krusell's hypothesis represents a Panglossian view of the world. If their explanation is correct, the negative welfare implications associated with many standard theories are completely reversed. Building the case for such a counterintuitive idea is difficult, but Hornstein and Krusell rise to the occasion. The paper builds, piece by piece, an accumulation of evidence that is consistent with their hypothesis. The result is a very interesting and novel paper that deserves serious attention.

The rest of my discussion will consist of three parts. I will first discuss the intuition of the hypothesis in the form of an analogy and then briefly review their evidence in favor of the hypothesis. Second, I will discuss other implications of the hypothesis, and show that they are not so favorable to the hypothesis. Finally, I will add some concluding remarks.

2. *Intuition and Evidence for the Hypothesis*

Anyone who has recently changed word-processing packages should find Hornstein and Krusell's arguments to be plausible. Changing word processors typically results in temporary decreases in true productivity. Simple tasks such as pagination and typing mathematical formulas suddenly become difficult. Adoption of the new technology temporarily decreases one's knowledge level, and requires effort in learning the new technology. Furthermore, standard measurements of output might not indicate any gain associated with adoption. A government statistician who measured output by page counts would not measure any change. Yet the *quality* of output has risen: the tables and equations are more readable and there might be increased compatibility with coauthors. Hornstein and Krusell argue that this experience is an important economic phenomenon.

In support of their argument, Hornstein and Krusell present four key pieces of evidence. I will briefly discuss each of these.

1. *The growth slowdown occurred in most developed countries.* Explanations for the productivity slowdown that are specific to the United States are less compelling because the slowdown was indeed an international phenomenon. Hornstein and Krusell's hypothesis passes the international test because technological progress tends to diffuse across developed countries. Thus, any increase in the rate of investment-specific technological progress should have been experienced internationally.
2. *The rate of decline of relative prices of producers' durable goods increased in the early to mid 1970s.* Figure 2 of the paper and the accompanying regression show that Gordon's measure of the relative price of durable goods declines more steeply after the early 1970s. If this decline is associated with an increase in the rate of technological progress (an interpretation that I will question later), then the data suggest that the timing is correct.
3. *The calibrated model produces slowdowns in measured productivity growth.* Both the learning effects and hypothesized quality measurement problems lead to simulated productivity slowdowns. To obtain decreases in *output growth,* however, the model requires either a very large compatibility problem or an exogenous increase in the quality content of output. The key parameters underlying both of these features are difficult to observe in practice. Hornstein and Krusell are only able to offer some suggestive evidence on the qualitative, but not quantitative, values of the parameters.
4. *Total factor productivity growth recovered almost completely in the measurable sector after 1979.* Hornstein and Krusell show that the productivity slowdown in the sectors in which quality mismeasurement is less problematic was indeed temporary. This fact is consistent with their hypothesis on learning leading to temporary slowdowns in true productivity growth, and mismeasurement leading to prolonged slowdowns in measured productivity growth.

3. *Some Opposing Evidence*

While Hornstein and Krusell's evidence favors their hypothesis, it should be noted that they have explored only a small subset of its potential implications. In this section, I will argue that there are other direct implications of the theory. I will then show that the data are not consistent with these implications, or with Hornstein and Krusell's interpretation of the underlying facts. Instead, the data on which I focus are more consistent with standard explanations for the productivity slowdown that focus on negative economic events.

In the spirit of the standard growth accounting literature, Hornstein and Krusell analyze only production functions, and do not specify the optimizing problem underlying technology adoption or the rest of the features of the general equilibrium economy. Conducting a complete analysis of a fully specified general equilibrium model is, of course, very difficult. Some of the general equilibrium principles and implications of technological change and quality mismeasurement, however, are quite intuitive, and I will focus on these elements.

I begin by emphasizing two principles about technological change and one about quality mismeasurement. None of these should be controversial. First, an increase in the rate of technological change implies that the economy's production possibilities frontier is expanding at a faster rate. Thus, an increase in technological change cannot make agents worse off in present-value terms. Second, firms are not *forced* to adopt new technology; instead, they can make an optimizing choice. A firm should only invest in new technology if the expected present discounted value of future profits outweighs current losses due to learning. The third principle regards quality mismeasurement. True increases in quality, although difficult for economists to measure, must be perceived by individuals, or else they are not actual increases in quality. (The only exceptions to this principle are "postexperience goods," such as drugs, whose true quality may not be perceived even with consumption.)

These three principles imply that even in a world with learning effects and difficulties in measuring quality, indices of economic performance that summarize individuals' information should accurately reflect the net positive impact of technological change. One such index is the stock-market value of firms. Although economists cannot accurately measure increases in production opportunities or increases in quality, firms, consumers, and shareholders should perceive these changes. If a positive shock to the growth rate of technology was an important part of the economic events of the 1970s, then the stock market should have reflected it.

Figure 1 LOG REAL STOCK PRICE (VALUE-WEIGHTED CRSP INDEX, DIVIDENDS INCLUDED)

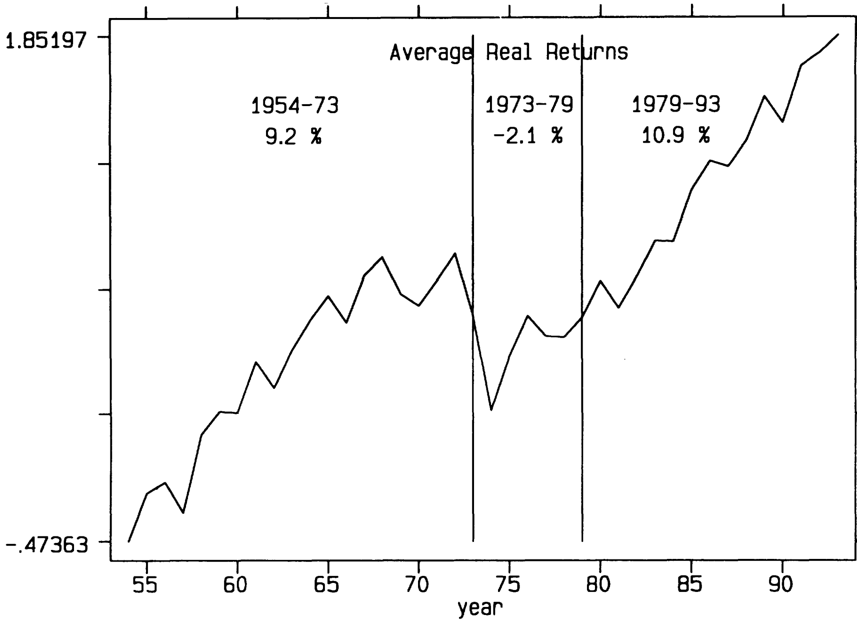

Figure 1 shows the behavior of the log of real stock prices (with dividends) from 1954 to 1993, as well as the average real returns. The CPI was used to deflate nominal prices; alternative indices led to similar results. The graph clearly shows that agents perceived the net effect of shocks in the early 1970s to be *negative:* the stock market fell during the 1970s, and real returns were negative from 1973 to 1979. Note that short-run learning effects cannot be invoked as an explanation, because the stock price should reflect the present discounted value of the effects of technological change. Furthermore, arguing that the oil shocks can explain the decline in the stock market supports the notion that the productivity slowdown was more likely due to negative economic shocks than to positive economic shocks. Finally, while one might try to explain an individual firm's decline in stock price with the argument of capital losses on existing equipment, the first principle I set out above suggests that this argument cannot explain an *aggregate* decline in the value of firms.

The evidence from stock prices does not refute the possibility that the rate of technological change increased in the early 1970s. It does suggest, however, that the negative shocks during this period were much more important than any positive shocks. Thus, the stock price evidence sup-

Figure 2 CORPORATE-PROFITS–GDP RATIO, NONELECTRICAL MACHINERY RELATIVE TO NONDURABLE MANUFACTURING

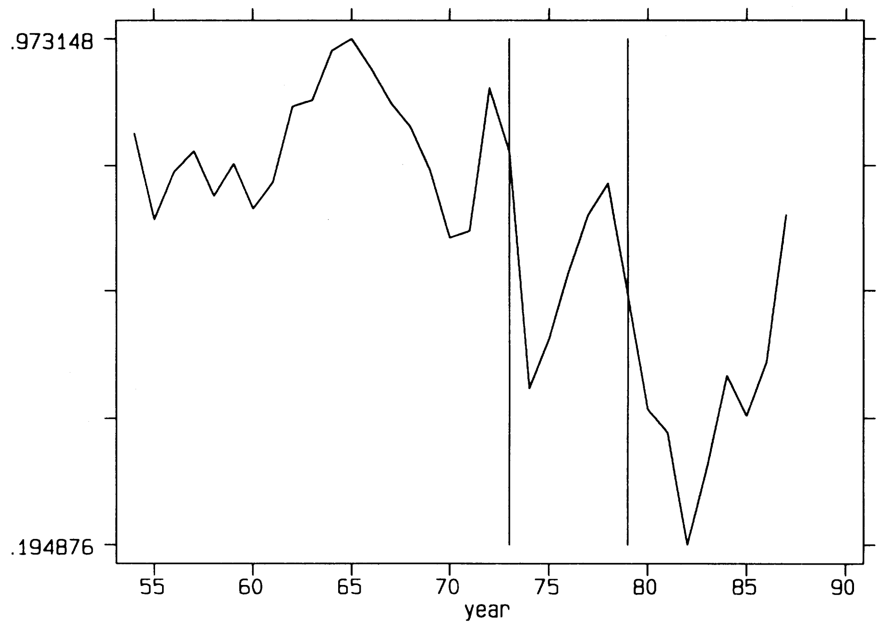

ports "negative shock," rather than "positive shock," explanations of the productivity slowdown.

I will suggest further that even the evidence for an increase in the rate of investment-specific technological change in the early 1970s is weak. The only evidence presented by Hornstein and Krusell in this regard is the decline in the relative price of producers' durable goods. If the source of this decline in relative prices is high relative productivity growth in producers' durable goods industries, then relative profit rates in these industries should have risen (or at least not fallen). It is difficult to imagine a scenario in which technological growth would lower relative profit rates. Yet relative profit rates of equipment producers fell. Figure 2 shows relative profit rates for nonelectrical machinery (SIC 35), constructed using data on corporate profits and GDP from the *Economic Report of the President*. This industry represents an important fraction of equipment manufacturing. Relative profits show downward trends from the mid-1960s to the early 1980s. The same is true for durable goods overall. These trends are at odds with the notion of high relative technological progress in producers' durable goods.

If not technological progress, then what caused the relative price of producers' durable goods to fall over this period? One possible explanation is a decline in markups. Many of these industries have historically been characterized by firms with market power and high rents. For example, the four-firm concentration ratio for turbines (SIC 3511) was 90% in 1972. One force that could have led to an erosion of that market power and a decline in markups is the rise of import competition. The import share of the domestic market for nonelectrical machinery rose from 3% in 1965 to 10.6% in 1980. Another possible explanation is that firms in these industries compete through quality improvements rather than through price. This strategy could explain both the increased emphasis on quality and the decline in the quality-adjusted price over this period. Thus, there is an alternative story that explains some of the facts highlighted by the authors. Completely independently of the source of the productivity slowdown, markups on producers' durable goods fell, leading to downward trends in relative prices and increased investment in equipment.

4. *Concluding Remarks*

Based on the evidence reviewed in the previous section, I am skeptical of the potential for increased rates of investment-specific technological change to explain the slowdown in both output and productivity growth that began in the early 1970s. The decline in relative profit rates in producers' durable goods calls into question the interpretation of price declines as evidence of technological change. The poor performance of the stock market in the 1970s supports a negative-shock rather than positive-shock explanation of the productivity slowdown.

On the other hand, the evidence is more favorable if one shifts the timing to explaining events of the 1980s and 1990s. As shown in Figure 1, stock-market returns were high during the period 1979–1993. Furthermore, despite heavy import competition, relative profit rates in producer's durable goods began to rise in the early 1980s. Thus, Hornstein and Krusell's evidence as well as these series shows patterns consistent with an increase in investment-specific technological change beginning in the 1980s. The recovery of productivity growth in the measurable sector during this period suggests that the negative impact of learning effects was small.

In sum, Hornstein and Krusell have presented a very interesting economic theory that links increased rates of investment-specific technological progress with productivity slowdowns. While the link exists in theory, the evidence does not support its application to the 1970s slowdown. The evidence is more favorable for several aspects of the hypothesis, such as increased technological progress and mismeasurement problems, if they are used to explain growth during the 1980s and 1990s.

Discussion

The authors began the discussion by replying to some issues raised in the formal comments. Per Krusell agreed that it was difficult to pin down the size and timing of the hypothesized increase in the rate of technological progress; but he felt that the weight of the numbers presented in their paper was in favor of such an increase having occurred. One interesting direction, Krusell suggested, would be to assess the sensitivity of the finding of increased technical progress to the choice of break date, in particular, to consider break dates later than 1973. Krusell also defended their model of quality change, in which the "share" of quality in final output is exogenous, as a reasonable approximation to a more fully specified model with endogenous quality change. He agreed that the development of endogenous-quality models, which might well look something like recent endogenous-growth models, was a worthwhile direction for research. Andreas Hornstein suggested that, while it was important to model the technology adoption choice in an explicit optimizing framework, the models studied in their paper provided a useful starting point for assessing the effects of increased technical progress on measured productivity. He also argued that their assumption that new-technology adoption has a large negative initial impact on productivity is not inconsistent with evidence from the growth and learning-by-doing literatures.

There was some discussion of Valerie Ramey's argument, that the poor stock-market performance of the 1970s seemed inconsistent with the premise of an increased rate of technological change during that decade. Ben Bernanke pointed out that the decade after 1973 included two oil shocks and two serious recessions, which might explain low returns to stocks despite underlying improvement in technology. He noted that more recently price–earnings ratios in the stock market have risen significantly.

Sam Kortum elaborated on Bob Gordon's suggestion, that the productivity slowdown was in fact a return to "normal" rates of technical progress, following an era that was truly exceptional in terms of economically important new technologies. He pointed out that this thesis, plus the assumption that innovations gradually diffuse, could account for the worldwide nature of the slowdown; it could also account for the more severe slowdown experienced by industrial countries other than the United States, which, having reached the technological frontier, could no longer enjoy the "catch-up" productivity growth bonus of technological followers. Gordon cited evidence that the U.S. output–capital ratio is declining as support for the view that diminishing returns to factors have begun to overwhelm technological improvement, as well as for the position that the productivity slowdown has occurred outside the investment sector.

Julio Rotemberg raised the puzzle that there has been a lot of investment in computers even though the measured returns on these investments have been very low. He thought that this might actually be consistent with the authors' argument, because these low rates of return could be due to learning. Moreover, rational firms might be investing in spite of these low rates of return because they expect high returns to accrue in the future. Krusell questioned standard growth accounting exercises that conclude that, because the share of information technology in investment is relatively low and the measured returns small, the new technologies cannot be having an important effect on productivity. He argued that electronic devices are becoming ubiquitous and affect the productivity, both measured and unmeasured, of many forms of capital, especially equipment.